arXiv:1404.5684v1 [math.NA] 23 Apr 2014

Compression Approaches for the Regularized Solutions of Linear Systems from Large-Scale Inverse Problems

Sergey Voronin[1,2], Guust Nolet[1], and Dylan Mikesell[3]

[1]Géoazur, Université de Nice, 06560 Sophia Antipolis, France
[2]Department of Applied Mathematics, University of Colorado, Boulder, CO 80309, USA
[3]Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

April 24, 2014

Abstract

We introduce and compare new compression approaches for obtaining regularized solutions of large linear systems which are commonly encountered in large scale inverse problems. We first describe how to approximate matrix vector operations with a large matrix through a matrix of smaller size, by borrowing from ideas used in wavelet image compression. Next, we describe and compare approaches based on the use of the low rank SVD factorization, which can result in further size reductions. We show some analytical results concerning the different methods and illustrate the application of the different approaches using a very large linear system from a Geotomography application, where we obtain significant compression gains with our methods, while still resolving the main features of the solutions.

1 Introduction

This paper introduces several approaches to obtain approximate but accurate regularized solutions to large linear systems arising from large scale inverse problems, without having to load into memory the often very large original matrix used in the solution to the corresponding optimization problems. Typically, such as in the case of the Geotomography application which we mention here for illustration [17] (involving the reconstruction of seismic wave velocities in the Earth's interior with respect to an existing spherically symmetric model), the physics calls for a solution of a linear system $Ax = \bar{b}$ with $A \in \mathbb{R}^{m \times n}$ (often with $m \neq n$) and where instead of the true right hand side $\bar{b}$ we are given the noisy approximation $b = \bar{b} + \nu$ with ν being the unknown noise vector. The matrix A can be very large and is likely to be ill-conditioned [14]. In order to obtain the solution given matrix A and right hand side b we often use some derivative of Tikhonov regularization [19]:

$$\bar{x} = \arg\min_x \left\{ ||Ax - b||_2^2 + \lambda ||x||_2^2 \right\}, \tag{1.1}$$

which is equivalent to constraining the ℓ_2 norm of the solution, depending on the value of λ. For large λ, $\bar{x}$ tends to be close to zero. Regularization is necessary to counter the effects of ill-conditioning: the presence of small singular values in the matrix, which if left unaccounted for, blows up the norm of the solution and makes it very sensitive to any data errors [2]. This property is worth repeating as it is central to the ideas in this paper: small errors in the operator A and the right hand side b do not induce big changes in the regularized solution. The regularization in (1.1) is referred to as ℓ_2 regularization, because it involves the minimization of the ℓ_2 norm. Other types of regularization are possible: for example, sparsity constrained regularization can also be used for Geophysical inverse problems [3]. In this paper, we discuss the application of our methods to ℓ_2 regularization, as it is the most commonly used regularization. However, the techniques apply also to other types of regularization and optimization techniques. The quadratic functional in (1.1) can be differentiated to yield the linear system for the regularized solution:

$$(A^T A + \lambda I)\bar{x} = A^T b. \tag{1.2}$$

If the matrix A is not too large, then there is no problem in solving this linear system with an iterative algorithm. A conjugate gradient or the LSQR algorithm [15] can be efficiently used for this purpose. Typically, we may wish to incorporate additional terms into the regularization, such as Laplacian smoothing [14]. In that case we solve instead:

$$\bar{x} = \arg\min_x \left\{ ||Ax - b||_2^2 + \lambda_1 ||x||_2^2 + \lambda_2 ||Lx||_2^2 \right\}, \tag{1.3}$$

which can be solved through the linear system:

$$(A^T A + \lambda_1 I + \lambda_2 L^T L)\bar{x} = A^T b, \tag{1.4}$$

or through the augmented least squares problem and its corresponding normal equations:

$$\bar{x} = \arg\min_x \left\| \begin{bmatrix} A \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix} x - \begin{bmatrix} b \\ 0 \\ 0 \end{bmatrix} \right\|_2^2 \implies \begin{bmatrix} A \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix}^T \begin{bmatrix} A \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix} \bar{x} = \begin{bmatrix} A \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix}^T \begin{bmatrix} b \\ 0 \\ 0 \end{bmatrix}.$$

As long as A and L can be applied to vectors, the solution can be obtained by a number of iterative algorithms. The problem occurs when A is too large to load into memory. In the Geotomography application we refer to [7, 17, 21], the matrix is several terabytes in size, so it may not be possible to load into memory in full, even on relatively large memory computer clusters. Thus, we must find ways to condense the matrix size using acceptable approximations which do not significantly alter the final regularized solutions.

Many attempts at approximating matrices have been documented [11, 22]. However, there have been few papers which attempt to apply the approximations to regularization. One of the main papers which precedes ours is [10], where Krylov subspace approximations for Tikhonov regularization are discussed. In this paper, we discuss two different techniques: wavelet based approximations and low rank SVD (singular value decomposition). Our SVD techniques are especially effective when the matrix is ill-conditioned. From our experiments, Krylov subspace techniques, while interesting and promising,

tend to do worse when the matrix is ill-conditioned in contrast to the techniques we describe, which do not significantly degrade the solution quality, while lowering the hardware requirements to obtain such solutions. Even if A is small enough that it can be loaded into memory, there may still be interest in the techniques we describe for gains of speed or to be able to solve several problems at once on one machine.

2 Approximate Matrix-Vector Operations with Wavelet Compression

Most iterative algorithms applicable to our discussion can be successfully implemented if we can perform the two key operations with the matrix A:

$$Ax \quad \text{and} \quad A^T y, \tag{2.1}$$

where for $A \in \mathbb{R}^{m\times n}$, $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$. We now discuss a technique to do these operations approximately, using a smaller matrix derived from A by means of wavelet compression [5, 9]. Wavelets provide a multiresolution approach to signal analysis, capturing long and short term wavelength parts of a signal, and wavelet based transformations can be performed efficiently [1, 18]. In our application, our matrix rows have features which are well represented by wavelets. To motivate this approach, consider wavelet compression applied to a synthetic Geophysical model (or any typical vectorized image). We compare the original model x (in row vector form) to the inverse transform of the thresholded wavelet transformed model based on the relation:

$$x \approx \left(W^{-1}\left(\mathrm{Thr}(Wx^T)\right)\right)^T,$$

where W and W^{-1} represent the forward and inverse wavelet transforms [13] and the thresholding operation $\mathrm{Thr}(\cdot)$ retains a certain percentage of the largest coefficients (by absolute value) of its input vector. The transpose operations assure that we are applying the transforms to column vectors, in view of their representation as matrices W and W^{-1}. The above relation assumes that row vector x is wavelet compressible. This is not necessarily the case for arbitrary x. However, in the case of the application we consider here, the vectors are Geophysical kernels representing a sensitivity of the observable (usually a phase or a delay) with respect to the intrinsic velocity as a function of space [12]. These kernels arise from integral equations and are generally smooth, and have been observed by us to be compressible by imposing a threshold on the wavelet coefficients. Many different kinds of thresholding functions exist. For our purposes, we simply use the hard thresholding function:

$$H_\alpha(x) = \begin{cases} x & \text{if } |x| > \alpha, \\ 0 & \text{if } |x| \le \alpha. \end{cases}$$

With the right choice of wavelet transform, only a small fraction of the coefficients in the wavelet transformed representation Wx need to be retained for a good reconstruction. That is, the threshold α can be taken to be quite large relative to the magnitudes of the elements of the vector Wx. In Figure 1, below, a smooth CDF transform was used [4].

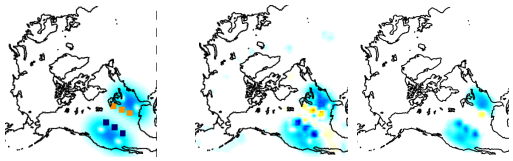

Figure 1: Synthetically constructed model and its representation with 15% and 3% of the largest wavelet coefficients.

Above, we compare the original row vectorized image x to the reconstructed image $\left(W^{-1}\,\mathrm{Thr}(Wx^T)\right)^T$ with fewer nonzero wavelet coefficients and observe that as the amount of nonzero wavelet coefficients retained is decreased, the reconstruction quality decreases, but the main features of the image are still retained. Clearly, the reconstruction error can be controlled by keeping a certain amount of nonzero coefficients.

We would like to apply the same principle to approximate matrix vector operations with the big original matrix A through a smaller matrix M so that only the smaller matrix M needs to be loaded into memory. The matrix M will have the same dimensions as A but fewer nonzeros, so it takes less space on disk and in memory. One forms this matrix by transforming and thresholding the individual rows of A. That is, each row of M below is obtained by applying the wavelet transform and thresholding to the corresponding row of A. The operations are visualized below:

$$A = \begin{bmatrix} r_1 \\ r_2 \\ \vdots \\ r_m \end{bmatrix} \;\rightarrow\; M = \begin{bmatrix} \mathrm{Thr}(Wr_1^T)^T \\ \mathrm{Thr}(Wr_2^T)^T \\ \vdots \\ \mathrm{Thr}(Wr_m^T)^T \end{bmatrix} = \mathrm{Thr}(AW^T) \approx AW^T$$

We can then approximate the operations (2.1). Using the relations:

$$Mx \approx AW^T x \quad \text{and} \quad M^T y \approx (AW^T)^T y = WA^T y,$$

we obtain the approximation formulas:

$$Ax \approx MW^{-T} x \quad \text{and} \quad A^T y \approx W^{-1} M^T y. \tag{2.2}$$

This means that the operations (2.1) can be performed in terms of the smaller matrix M and the inverse and inverse-transpose wavelet transforms. In practice, only M needs to be loaded in memory. The inverse-transpose transform is equivalent to the forward transform when W is orthogonal and $W^{-1} = W^T$. For the non-orthogonal case, such as for example the CDF 9/7 transform, the inverse-transpose transform can be programmed by applying the forward transform with the inverse filters. The success of this approximation method depends on the size ratio between M and A and the percent error in the approximate operations. This depends on the data, the transform that is used, and the threshold used in the thresholding function. Typically, we identify the threshold as follows. The input is sorted by putting the entries with largest absolute magnitude in front. Then a threshold is identified by putting the marker at some point of the nonzero entries (for example at the largest 15% mark of the total nonzeros). Then all the entries with absolute magnitude less than the identified threshold

are zeroed out. The percent error in the approximate operations then depends on the percent error in the reconstruction of each row. That is, if for an arbitrary row r, $\left(W^{-1}\,\mathrm{Thr}(Wr^T)\right)^T$ is not close to r, then the approximate operations using M formed with this threshold will probably not be accurate. A less aggressive threshold then needs to be used. Later we give examples for our Geotomography application. For our data set from the Geotomography application, we have observed that one can expect M to be at least 3 times smaller in memory requirements than A without incurring errors greater than about 10 percent in the operations Ax, A^Ty, and A^TAx. Notice that if one plugs in the compressed wavelet approximations for the matrix operations into (1.2), one gets:

$$(W^{-1}M^TMW^{-T} + \lambda I)\tilde{x} = W^{-1}M^Tb.$$

If A is too big to load directly into memory, the first approach one may think of is to split the matrix in parts along its rows, with the matrix vector operations applied blockwise:

$$A = \begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_p \end{bmatrix} \implies Ax = \begin{bmatrix} A_1x \\ A_2x \\ \vdots \\ A_px \end{bmatrix} \quad \text{and} \quad A^Ty = \begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_p \end{bmatrix}^T \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_p \end{bmatrix} = \sum_{j=1}^{p} A_j^T y_j.$$

In theory, if hard disks were very fast, we could keep only one submatrix in memory at a time to do these operations, but since disk read and write times for large matrices are usually quite slow, this is not a feasible solution. Instead, we can apply the wavelet compressed technique to the block matrices. We can proceed to form the matrices $M_1 = \mathrm{Thr}(A_1W_1^T), \ldots, M_p = \mathrm{Thr}(A_pW_p^T)$, which are smaller wavelet thresholded versions of the original blocks. We can then perform approximate operations using these new sparser blocks:

$$A = \begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_p \end{bmatrix} \to M = \begin{bmatrix} \mathrm{Thr}(A_1W_1^T) \\ \mathrm{Thr}(A_2W_2^T) \\ \vdots \\ \mathrm{Thr}(A_pW_p^T) \end{bmatrix} \implies Ax \approx \begin{bmatrix} M_1W_1^{-T}x \\ M_2W_2^{-T}x \\ \vdots \\ M_pW_p^{-T}x \end{bmatrix} \quad \text{and} \quad A^Ty \approx \sum_{j=1}^{p} W_j^{-1}M_j^Ty_j. \tag{2.3}$$

In the above formulas, we have used different transform matrices $W_1, \ldots, W_p$ for the different blocks. In practice, one would often use $W_1 = \cdots = W_p = W$ but it is possible to use different transforms on different blocks if necessary. Note here that M can be formed through the blocks $M_1, \ldots, M_p$ no matter how large A is. We can always choose p large enough so that the individual blocks A_j are manageable in size. Of course, it is possible that after forming M it may still be too large to be loaded into memory as a whole. In that case, we can still do operations in blocked form (2.3) by loading as many parts of M as we can into memory, performing part of the operation and then replacing the in-memory blocks with the remaining blocks of M to perform the rest. As long as fast disks (such as SSDs) are available, this is viable in practice, but may be very slow if many operations are needed.

3 Low Rank SVD

The wavelet approximation techniques for matrix-vector operations discussed in the previous section enable us to approximate the operations (2.1) through a matrix several times smaller than A. However, in practice, the matrix M can still be quite big if A is particularly large. It is plausible that we can do some operations with A through M but only for a relatively short amount of time (perhaps through the blocked form 2.3). Assuming that we can indeed do a limited number of matrix vector multiplications with A through M, we now discuss other techniques for compression based on the low rank singular value decomposition (SVD). Once such a decomposition is obtained through a limited amount of matrix vector multiplications with A (approximated through M), we can obtain approximate forms of regularization algorithms which require the use of significantly smaller matrices.

3.1 Overview

Any matrix $A \in \mathbb{R}^{m \times n}$ has the SVD decomposition [20] such that:

$$A = \sum_{i=1}^{\min(m,n)} \sigma_i u_i v_i^T = \sum_{i=1}^{r} \sigma_i u_i v_i^T = U \Sigma V^T \ , \ \Sigma = \text{diag}(\sigma_1, \ldots, \sigma_r, 0, \ldots, 0) \ , \ \sigma_i \in \mathbb{R}, u_i \in \mathbb{R}^m, v_i \in \mathbb{R}^n,$$

where U and V are orthogonal matrices ($U^T U = U U^T = V^T V = V V^T = I$) with $\sigma_j = 0$ for $j > r$, where $r = \text{rank}(A) \leq \min(m, n)$. The sizes of the matrices are: U is $m \times m$, Σ is $m \times n$ and V is $n \times n$. There are r nonzero singular values: $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > 0$. For a matrix which is not well conditioned, many nonzero singular values σ_j will be very small relative to the largest singular value σ_1 and the drop off in value starting from σ_1 will be rapid and nonlinear. Consider the low rank approximation of A by taking into account only the first k singular values and vectors:

$$A_k = \sum_{i=1}^{k} \sigma_i u_i v_i^T = U_k \Sigma_k V_k^T \approx U \Sigma V^T,$$

with $k < r$, $U_k \in \mathbb{R}^{m \times k}$ consisting of the first k columns of U, $\Sigma_k = \text{Diag}(\sigma_1, \ldots, \sigma_k) \in \mathbb{R}^{k \times k}$, and $V_k \in \mathbb{R}^{n \times k}$ consisting of the first k columns of V. When $k \ll r \leq \min(m, n)$, these matrices are significantly smaller than the corresponding full SVD matrices U, Σ, and V. When the matrix A is of low numeric rank (few singular values are above machine precision relative to matrix dimensions), this approximation can be accurate even for small k. In any case, the matrix A_k is the best rank k approximation of A with the bound:

$$||Az - A_k z||_2 \leq \sigma_{k+1} ||z||_2 \quad \forall z \in \mathbb{R}^n.$$

The choice of k is up to the user, but greater k requires greater computation time and storage requirements. Notice that for $k < r$ we have the following expressions:

$$\begin{aligned}
U &= [U_k, \hat{U}_k] \quad ; \quad V = [V_k, \hat{V}_k]; \\
A &= \sum_{i=1}^{k} \sigma_i u_i v_i^T + \sum_{i=k+1}^{r} \sigma_i u_i v_i^T = U_k \Sigma_k V_k^T + \hat{U}_k \hat{\Sigma}_k \hat{V}_k^T = A_k + \hat{U}_k \hat{\Sigma}_k \hat{V}_k^T = A_k + \hat{A}_k, \\
A^T &= \sum_{i=1}^{k} \sigma_i v_i u_i^T + \sum_{i=k+1}^{r} \sigma_i v_i u_i^T = V_k \Sigma_k U_k^T + \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T = A_k^T + \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T = A_k^T + \hat{A}_k{}^T, \\
A^T A &= \sum_{i=1}^{k} \sigma_i^2 v_i v_i^T + \sum_{i=k+1}^{r} \sigma_i^2 v_i v_i^T = V_k \Sigma_k^2 V_k^T + \hat{V}_k \hat{\Sigma}_k^2 \hat{V}_k^T = A_k^T A_k + \hat{V}_k \hat{\Sigma}_k^2 \hat{V}_k^T = A_k^T A_k + \hat{A}_k{}^T \hat{A}_k,
\end{aligned}$$

where

$$A_k = \sum_{i=1}^{k} \sigma_i u_i v_i^T = U_k \Sigma_k V_k^T \quad \text{and} \quad A_k^T A_k = \sum_{i=1}^{k} \sigma_i^2 v_i v_i^T = V_k \Sigma_k^2 V_k^T,$$

and $U_k^T \hat{U}_k = V_k^T \hat{V}_k = 0$ and $U_k^T U_k = \hat{U}_k^T \hat{U}_k = V_k^T V_k = \hat{V}_k^T \hat{V}_k = I$. Additionally, we have the following properties:

Lemma 3.1 *For vectors $u_1 \in \mathbb{R}^k$ and $u_2 \in \mathbb{R}^m$, $||U_k u_1||_2 = ||u_1||_2$ and $||U_k^T u_2||_2 \leq ||u_2||_2$. And the same for vectors $v_1 \in \mathbb{R}^k$ and $v_2 \in \mathbb{R}^n$ and V_k and V_k^T.*

Proof. Note that

$$UU^T = I = [U_k, \hat{U}_k] \begin{bmatrix} U_k^T \\ \hat{U}_k^T \end{bmatrix} = U_k U_k^T + \hat{U}_k \hat{U}_k^T \implies U_k U_k^T = I - \hat{U}_k \hat{U}_k^T.$$

Thus:

$$\begin{aligned}
&||U_k u_1||_2^2 = \langle U_k u_1, U_k u_1 \rangle = \langle u_1, U_k^T U_k u_1 \rangle = \langle u_1, u_1 \rangle = ||u_1||_2^2, \\
&||U_k^T u_2||_2^2 = \langle U_k^T u_2, U_k^T u_2 \rangle = \langle u_2, U_k U_k^T u_2 \rangle = \langle u_2, (I - \hat{U}_k \hat{U}_k^T) u_2 \rangle = \langle u_2, u_2 \rangle - \langle u_2, \hat{U}_k \hat{U}_k^T u_2 \rangle \leq ||u_2||_2^2.
\end{aligned}$$

The computations with V_k and V_k^T take similar form. □

3.2 Computation with Randomized Algorithm

We now discuss how the low rank SVD may be computed. One direct way is to compute it from the full SVD of the matrix. Given the full SVD $A = U\Sigma V^T$ one can take the first k columns of U and V to be the matrices U_k and V_k and the first k diagonal elements of Σ to form Σ_k. For large matrices, this is not practical since the computation of the full SVD is prohibitively expensive (the cost for an $m \times n$ matrix is on the order of $\mathcal{O}(mn \min(m, n))$ operations [20]). The algorithm which we use is

an adaptation of the method proposed in [8]. The cost of the proposed randomized algorithm for the rank k SVD approximation is substantially lower (the cost is $\mathcal{O}(mnk)$ operations).

The randomized algorithm finding a rank k approximation of $A \in \mathbb{R}^{m\times n}$ proposed by [8] consists of several simple steps. The main idea is to obtain a good estimate for the range of A by forming products of A with a sample of random vectors, then using the orthogonal basis of this sample matrix to project the original matrix into a smaller, lower dimensional one, of which we extract the full SVD and use these components to construct the low rank SVD of the original big matrix A. The steps are as follows:

- Take k samples of the range of matrix A by multiplying A with random Gaussian vectors to form sample matrix Y of size $m \times k$. We then have $\operatorname{range} Y \approx \operatorname{range} A$.
- Orthogonalize this set of samples forming the matrix Q. Then $\operatorname{range} Q \approx \operatorname{range} A \implies QQ^TA \approx A$.
- Project the original matrix into a lower dimensional one: $B = Q^TA$ where B is $k\times n$, substantially smaller than A which is $m \times n$.
- Take the SVD of the smaller matrix $B = \tilde{U}_k\Sigma_kV_k^T$.
- Take as low rank SVD of A the product $U_k\Sigma_kV_k^T$ with $U_k = Q\tilde{U}_k$ (since $QQ^TA \approx A$).

Since the matrix B of size $k \times n$, can still be quite large for large n, we compute the SVD components $\tilde{U}_k$ and Σ_k of B using the eigendecomposition of the small $k\times k$ symmetric matrix BB^T and obtain V_k by applying B^T. This way, we avoid building B or taking the SVD of it directly. We use the following relations:

$$B = \tilde{U}_k\Sigma_kV_k^T = \sum_{i=1}^{k}\sigma_i\tilde{u}_iv_i^T \quad ; \quad B^T = V_k\Sigma_k\tilde{U}_k^T \quad ; \quad Bv_i = \sigma_i\tilde{u}_i;$$

$$BB^T = \left(\sum_{i=1}^{k}\sigma_i\tilde{u}_iv_i^T\right)\left(\sum_{j=1}^{k}\sigma_j\tilde{u}_jv_j^T\right)^T = \sum_{i,j=1}^{k}\sigma_i\sigma_j\tilde{u}_iv_i^Tv_j\tilde{u}_j^T = \sum_{i=1}^{k}\sigma_i^2\tilde{u}_i\tilde{u}_i^T = \tilde{U}_kD_k\tilde{U}_k^T.$$

This means the eigendecomposition of the $k \times k$ matrix BB^T gives us the low rank components $U_k = Q\tilde{U}_k$ and $\Sigma_k = \sqrt{D_k}$ element-wise. To compute the right eigenvectors v_i, we can use the following relations:

$$B^T\tilde{U}_k = V_k\Sigma_k\tilde{U}_k^T\tilde{U}_k = V_k\Sigma_k \implies B^T\tilde{U}_k\Sigma_k^{-1} = V_k,$$

which implies:

$$v_i = V_ke_i = (B^T\tilde{U}_k\Sigma_k^{-1})e_i = \frac{1}{\sigma_i}B^T\tilde{u}_i = \frac{1}{\sigma_i}A^TQ\tilde{u}_i, \tag{3.1}$$

assuming all the singular values in Σ_k are above zero (which is the case for k smaller than the numerical rank).

Notice that all matrix-vector operations involving A and A^T can be approximated via the wavelet compressed matrices M and M^T. To build up BB^T column by column we can use matrix-vector products with standard basis vectors e_j:

$$BB^T e_j = Q^T A A^T Q e_j \approx Q^T M W^{-T} W^{-1} M^T Q e_j, \tag{3.2}$$

and for the right eigenvectors, we have from (3.1) that:

$$v_i = \frac{1}{\sigma_i} A^T Q u_i \approx \frac{1}{\sigma_i} W^{-1} M^T Q u_i.$$

We now illustrate the main steps of the random algorithm to compute the low rank SVD, which we use in our computations for the numerical experiments. Below, we use Matlab like pseudocode.

- Take k samples of matrix A with random Gaussian vectors and perform Gram-Schmidt orthogonalization to calculate the projection matrix Q.

```
for j=1:k
    rj = randn(n,1);
    yj = A*rj;
    Y(:,j) = yj;
end

Q = Y;
for ind=1:2
    for j=1:k
        vj = Q(:,j);
        for i=1:(j-1)
            vi = Q(:,i);
            vj = vj - project_vec(vj,vi);
        end
        vj = vj/norm(vj);
        Q(:,j) = vj;
    end
end
```

  where the projection of v in direction of u is defined as $\frac{(v\cdot u)}{||u||_2^2}u$. For best results, the Gram-Schmidt orthogonalization should be performed twice to account for loss of orthogonality. Note that for matrix-vector multiplications with A we use $Ar_k \approx MW^{-T}r_k$.

- Build the $k \times k$ matrix $BB^T = Q^T A A^T Q$ by computing k matrix-vector products with standard basis vectors.

  Once we have built Q and its transpose, we can form the matrix BB^T column by column:

```
BBt = zeros(k,k);
for j=1:k
```

```
    ej = zeros(k,1);
    ej(j) = 1;
    colj = Qt*(A*(At*(Q*ej)));
    BBt(:,j) = colj;
end
```

Here, we would make use of (3.2) for approximating $Q^T A A^T Q e_j$.

- Compute the eigendecomposition of BB^T

  This is just the eigendecomposition of a small $k \times k$ matrix:

```
[Uhat,D] = eig(BBt);
```

- Compute the low rank SVD components of A by using the eigendecomposition derived in the previous step and applying $B^T = A^T Q$ to eigenvectors.

  Here we use the fact that the eigenvalues of BB^T are the squares of the singular values of B and the computation (3.1) for the eigenvectors V.

```
Sigma = zeros(k,k);
for i=1:k
    Sigma(i,i) = sqrt(D(i,i));
end

U = Q * Uhat;

V = zeros(n,k);
for j=1:k
    vj = 1/Sigma(j,j) * (At * U(:,j));
    V(:,j) = vj;
end
```

  where we would use $A^T u_j \approx W^{-1} M^T u_j$.

We note that the implementation of the low rank SVD algorithm above is simple, as long as we can perform matrix-vector operations using the wavelet compressed matrix M and compute the eigendecomposition of a small $k \times k$ matrix, which can be done with a large number of available numerical packages. The disadvantage of this version is that working with the matrix BB^T essentially squares the condition number of A, such that small singular values near machine precision may not be properly resolved. This is an issue if A is expected to have very small singular values amongst $\sigma_1, \ldots, \sigma_k$. However, if we take k to be small relative to $\min(m, n)$ as we do in our application, σ_k is significantly larger in magnitude than machine precision. The implementation of the algorithm in the pseudocode above is not very efficient for the randomized algorithm proposed in [8], but one that is practical to use for very large A when the corresponding wavelet compressed matrix $M = \mathrm{Thr}(AW^T)$ is available. In particular, for a more efficient implementation, one may want to block as many operations as possible, replacing matrix-vector by matrix-matrix multiplications. If possible, one may want to explicitly

compute the matrix B and then use it to form BB^T. Likewise, V_k can be calculated directly from the matrix product $B^T\tilde{U}_k\Sigma_k^{-1}$.

3.3 Application to Regularization Schemes

For purposes of iterative regularization algorithms, we can make use of the low rank SVD in several ways. If we obtain the low rank SVD of the whole matrix, we can directly use it to approximate matrix vector operations:

$$Ax \approx U_k\left(\Sigma_k(V_k^T x)\right) \quad \text{and} \quad A^T y \approx V_k\left(\Sigma_k(U_k^T y)\right), \tag{3.3}$$

and in some situations this is the most convenient and straightforward approach. The disadvantage of this approach is that one must keep the matrices U_k, U_k^T, V_k, V_k^T in memory. Here and below we do not pay attention to storing the matrix Σ_k which is a very small diagonal matrix. If the matrix A is large it may be difficult to compute the low rank SVD of the whole matrix A. Instead, if we block A as previously discussed, we can compute the low rank SVD of certain blocks or of each block. In some applications, it may be possible to arrange the blocks of A in a way that the first block of A contains many linearly dependent rows. If that is the case, then it is worthwhile to take the low rank SVD of that block since it could be approximated well with small k. We can then write down mixed relations as follows:

$$Ax \approx \begin{bmatrix} U_{k_1}\Sigma_{k_1}V_{k_1}^T x \\ M_2 W_2^{-T} x \\ \vdots \\ M_p W_p^{-T} x \end{bmatrix} \quad \text{and} \quad A^T y \approx V_{k_1}\Sigma_{k_1}U_{k_1}^T y_1 + \sum_{j=2}^{p} W_j^{-1} M_j^T y_j, \tag{3.4}$$

where in this example we have used the low rank SVD approximation for the first part of the matrix and the wavelet based approximation for the other parts.

Additional information can be learned by plugging in the low rank SVD directly into the regularization system. Our general model problem and it's corresponding linear system are:

$$\bar{x} = \arg\min_x \left(||Ax-b||_2^2 + \lambda_1||x||_2^2 + \lambda_2||Lx||_2^2\right) \implies (A^TA + \lambda_1 I + \lambda_2 L^T L)\bar{x} = A^T b. \tag{3.5}$$

Replacing all instances of A by the low rank SVD we have:

$$(A_k^T A_k + \lambda_1 I + \lambda_2 L^T L)\tilde{x_1} = A_k^T b,$$

which when expanded gives:

$$(V_k\Sigma_k^2 V_k^T + \lambda_1 I + \lambda_2 L^T L)\tilde{x_1} = V_k\Sigma_k U_k^T b. \tag{3.6}$$

The advantage of (3.6) is that if the right hand side $V_k\Sigma_k U_k^T b$ is computed at the start of the iteration, only the matrices V_k and V_k^T must be kept in memory during the iteration. We may think of precomputing the right hand side $A^T b$ and approximating only the operator $A^T A$. Note that $A^T b$ can always be precomputed before the iteration as long as we can split up A into blocks. In this case we get:

$$(V_k\Sigma_k^2 V_k^T + \lambda_1 I + \lambda_2 L^T L)\hat{x_1} = A^T b.$$

As we will show later, this can lead to slightly better results, though the norm of the solution for the same choice of λ_1 would be higher in this case. Another approach is to work with the lower dimensional projected system:

$$(U_k^T A)x = U_k^T b, \tag{3.7}$$

where $U_k^T A$ is $k \times n$ if $A \in \mathbb{R}^{m \times n}$. Note that we have the following simple result:

Lemma 3.2 *Given the low rank SVD $A_k = U_k \Sigma_k V_k^T$ of A, we have that $U_k^T A = U_k^T A_k = \Sigma_k V_k^T$.*

Proof. First, $U_k^T A_k = U_k^T (U_k \Sigma_k V_k^T) = \Sigma_k V_k^T$. Also:

$$U_k^T A = U_k^T \left(U_k \Sigma_k V_k^T + \hat{U}_k \hat{\Sigma}_k \hat{V}_k^T \right) = \Sigma_k V_k^T + 0 = \Sigma_k V_k^T.$$

□

If we solve (3.7) by means of Tikhonov regularization:

$$\tilde{x_2} = \arg\min_x \left\{ ||(U_k^T A)x - U_k^T b||_2^2 + \lambda_1 ||x||_2^2 + \lambda_2 ||Lx||_2^2 \right\} \tag{3.8}$$

$$\implies \quad \left((U_k^T A)^T (U_k^T A) + \lambda I + L^T L \right) \tilde{x_2} = (U_k^T A)^T U_k^T b, \tag{3.9}$$

we will obtain the same solution as (3.6):

Lemma 3.3 *The approximation scheme $(V_k \Sigma_k^2 V_k^T + \lambda_1 I + \lambda_2 L^T L)\tilde{x} = V_k \Sigma_k U_k^T b$ has the same solution as the Tikhonov regularized solution of the projected system $(U_k^T A)x = U_k^T b$.*

Proof. Since

$$\begin{aligned} & A = U_k \Sigma_k V_k^T + \hat{U}_k \hat{\Sigma}_k \hat{V}_k^T \implies U_k^T A = \Sigma_k V_k^T + 0 = \Sigma_k V_k^T \\ \implies \quad & (U_k^T A)^T (U_k^T A) = (U_k^T A_k)^T (U_k^T A_k) = (\Sigma_k V_k^T)^T (\Sigma_k V_k^T) = V_k \Sigma_k^2 V_k^T, \end{aligned}$$

the linear system is equivalent to:

$$\left(V_k \Sigma_k^2 V_k^T + \lambda_1 I + \lambda_2 L^T L \right) \tilde{x_2} = (U_k^T A)^T U_k^T b = A^T U_k U_k^T b.$$

Next, for the right hand side we have:

$$A^T U_k = V_k \Sigma_k U_k^T U_k + \hat{V}_k \hat{\Sigma}_k \hat{U_k^T} U_k = V_k \Sigma_k I + 0 = V_k \Sigma_k \implies A^T U_k U_k^T b = V_k \Sigma_k U_k^T b.$$

Hence the solution of (3.8) is equivalent to (3.6):

$$(V_k \Sigma_k^2 V_k^T + \lambda_1 I + \lambda_2 L^T L)\tilde{x_2} = V_k \Sigma_k U_k^T b.$$

□

The advantage of (3.7) is that it may be convenient for larger systems where we can only perform the low rank SVD of its blocks. In that case, we may form the blocked system:

$$\begin{bmatrix} U_{k_1}^T A_1 \\ U_{k_2}^T A_2 \\ \vdots \\ U_{k_p}^T A_p \end{bmatrix} x = \begin{bmatrix} U_{k_1}^T b_1 \\ U_{k_2}^T b_2 \\ \vdots \\ U_{k_p}^T b_p \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} \Sigma_{k_1} V_{k_1}^T \\ \Sigma_{k_2} V_{k_2}^T \\ \vdots \\ \Sigma_{k_p} V_{k_p}^T \end{bmatrix} x = \begin{bmatrix} U_{k_1}^T b_1 \\ U_{k_2}^T b_2 \\ \vdots \\ U_{k_p}^T b_p \end{bmatrix}, \tag{3.10}$$

and solve the optimization problem via the augmented normal equations:

$$\begin{bmatrix} U_{k_1}^T A_1 \\ U_{k_2}^T A_2 \\ \vdots \\ U_{k_p}^T A_p \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix}^T \begin{bmatrix} U_{k_1}^T A_1 \\ U_{k_2}^T A_2 \\ \vdots \\ U_{k_p}^T A_p \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix} \tilde{x_2} = \begin{bmatrix} U_{k_1}^T A_1 \\ U_{k_2}^T A_2 \\ \vdots \\ U_{k_p}^T A_p \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix}^T \begin{bmatrix} U_{k_1}^T b_1 \\ U_{k_2}^T b_2 \\ \vdots \\ U_{k_p}^T b_p \\ 0 \\ 0 \end{bmatrix} \quad \text{or} \quad \begin{bmatrix} \Sigma_{k_1} V_{k_1}^T \\ \Sigma_{k_2} V_{k_2}^T \\ \vdots \\ \Sigma_{k_p} V_{k_p}^T \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix}^T \begin{bmatrix} \Sigma_{k_1} V_{k_1}^T \\ \Sigma_{k_2} V_{k_2}^T \\ \vdots \\ \Sigma_{k_p} V_{k_p}^T \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix} \tilde{x_2} = \begin{bmatrix} \Sigma_{k_1} V_{k_1}^T \\ \Sigma_{k_2} V_{k_2}^T \\ \vdots \\ \Sigma_{k_p} V_{k_p}^T \\ \sqrt{\lambda_1} I \\ \sqrt{\lambda_2} L \end{bmatrix}^T \begin{bmatrix} U_{k_1}^T b_1 \\ U_{k_2}^T b_2 \\ \vdots \\ U_{k_p}^T b_p \\ 0 \\ 0 \end{bmatrix}.$$

The number of eigenvectors for each block can be adjusted based on their conditioning. If the same k is used for all the blocks then some are bound to be projected less accurately than others. If the right hand side is precomputed, only the matrices V_k^T and Σ_k must be in memory for each block. If it is easier to compute the eigenvector matrix U_k, then the default system with $U_k^T A$ may be useful.

A more aggressive approach is to use the right eigenvectors V_k to project the system from both sides to form a matrix of size $k \times k$. Instead of solving the full system:

$$(A^T A + \lambda_1 I + \lambda_2 L^T L)\bar{x} = A^T b,$$

we project the matrix used to a smaller space by multiplying on left by V_k^T and preconditioning on the right by V_k:

$$V_k^T (A^T A + \lambda_1 I + \lambda_2 L^T L)(V_k \tilde{y_3}) = V_k^T A^T b \quad ; \quad \tilde{x_3} = V_k \tilde{y_3}.$$

Expanding this and noting that $V_k^T V_k = I$, we have:

$$\left(V_k^T A^T A V_k + \lambda_1 I + \lambda_2 V_k^T L^T L V_k\right) \tilde{y_3} = V_k^T A^T b \quad ; \quad \tilde{x_3} = V_k \tilde{y_3}. \tag{3.11}$$

The key observation is that the matrix used in the linear system is $V_k^T A^T A V_k$, which is just of size $k \times k$, much smaller than the $m \times n$ matrix A. We can further simplify (3.11) using the following calculations:

Lemma 3.4 *Given the low rank SVD $A_k = U_k \Sigma_k V_k^T$ of A, we have that $V_k^T A^T A V_k = V_k^T A_k^T A_k V_k = \Sigma_k^2$ and $V_k A^T b = V_k A_k^T b = \Sigma_k U_k^T b$.*

Proof.

$$\begin{aligned} & V_k^T A^T = V_k^T (V_k \Sigma_k U_k^T + \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T) = V_k^T A_k^T = \Sigma_k U_k^T \implies A V_k = A_k V_k = (\Sigma_k U_k^T)^T = U_k \Sigma_k \\ \implies \quad & V_k^T A^T A V_k = \Sigma_k U_k^T U_k \Sigma_k = \Sigma_k^2 \\ \implies \quad & V_k A_k^T b = \Sigma_k U_k^T b. \end{aligned}$$

□

Thus, we can rewrite (3.11) as:

$$\left(\Sigma_k^2 + \lambda_1 I + \lambda_2 V_k^T L^T L V_k\right) \tilde{y_3} = \Sigma_k U_k^T b \quad ; \quad \tilde{x_3} = V_k \tilde{y_3}. \tag{3.12}$$

We will show later than when $\lambda_2 = 0$, $\tilde{x_3} = \tilde{x_1}$, an important result, since the system for $\tilde{y_3}$ can be solved on a small machine in that case. When $\lambda_2 \neq 0$, this is only an approximation. We can obtain the k columns of $V_k^T L^T L V_k$ by evaluating matrix vector products:

$$V_k^T L^T L V_k e_j \quad \text{for} \quad j = 1, \ldots, k.$$

This is feasible to do in practice, since k is not very large. This method is useful when many solutions with different values of λ_1 and λ_2 are desired.

Let us now summarize the different techniques we have described for approximate ℓ_2 regularization using the low rank SVD and their computational requirements.

(1) We can implement $(A^T A + \lambda_1 I + \lambda_2 L^T L)\bar{x} = A^T b$ as usual and replace the operations Ax and $A^T y$ with $U_k \Sigma_k V_k^T x$ and $V_k \Sigma_k U_k^T y$. This requires one to have the matrices U_k, U_k^T, V_k, V_k^T in memory, which may not be very efficient. However, this direct approach may be useful for larger matrices split into blocks using relations such as (3.4), where the low rank SVD is applied only to certain blocks and not to the whole matrix. In that case, only the SVD components for the relevant blocks need to be loaded.

(2) We can plug in the low rank SVD into the regularization problem to get the system:

$$(V_k \Sigma_k^2 V_k^T + \lambda_1 I + \lambda_2 L^T L)\tilde{x_1} = V_k \Sigma_k U_k^T b.$$

Note that the right hand side $V_k \Sigma_k U_k^T b$ can be precomputed before the iteration so that only the matrices V_k and V_k^T need to be in memory during iteration. The result should be equivalent to the first case but this approach is more efficient. Additionally, we can precompute accurately the right hand side $A^T b$ and use the system:

$$(V_k \Sigma_k^2 V_k^T + \lambda_1 I + \lambda_2 L^T L)\hat{x_1} = A^T b.$$

Here the only difference is in the right hand side. As we will see later this can sometimes lead to better solutions, but should be used with a larger threshold for λ_1.

(3) We can utilize the lower dimensional projected system $U_k^T A x = U_k^T b$. The corresponding system:

$$\left((U_k^T A)^T (U_k^T A) + \lambda_1 I + \lambda_2 L^T L\right) \tilde{x_2} = (U_k^T A) U_k^T b$$

is equivalent to the system for $\tilde{x_1}$. However, in certain cases the matrix U_k may be easier to compute than V_k (depending on the dimensions of $A^T A$ and $A A^T$) in which case one may then compute $U_k^T A$ by means of matrix-vector products $A^T U_k e_j$ for $j = 1, \ldots, k$. The method may also be useful for large systems since we can make use of (3.10).

(4) We can use the $k \times k$ system:

$$\left(\Sigma_k^2 + \lambda_1 I + \lambda_2 V_k^T L^T L V_k\right) \tilde{y_3} = \Sigma_k U_k^T b \quad ; \quad \tilde{x_3} = V_k \tilde{y_3}$$

The solution of the linear system can be done on small memory computers since it involves the use of $k \times k$ matrices only and one multiplication with V_k at the end. The last step can be performed on a larger machines loading only V_k into memory; or on smaller machines in blocks. This scheme is useful when many runs with the system with different values of λ_1 and λ_2 are desired. The solution is equivalent to $\tilde{x_1}$ when $\lambda_2 = 0$.

Note that up to now we have discussed the application of the compression techniques to ℓ_2 norm minimization problems. However, the techniques are applicable to other types of regularization also. For example, for ℓ_1 regularization, where we minimize $||x||_1$ instead of $||x||_2$, one typically uses a scheme similar to the Iterative Soft Thresholding Algorithm [6]:

$$x^{n+1} = \mathbb{S}_\tau \left(x^n + A^T b - A^T A x^n\right),$$

where $\left(\mathbb{S}_\tau(x)\right)_k = \text{sgn}(x_k) \max\{0, |x_k| - \tau\}$ is the componentwise soft thresholding function. The main computational requirement here is in the operation $A^T A x^n$, just as for ℓ_2 regularization. Hence, many of the techniques we have described can be used for different types of regularization.

3.4 Further Analysis and Error Bounds

In this section, we give more analysis for the SVD based schemes we have discussed. To make the analysis easier, we assume that $\lambda_1 = \lambda$ and $\lambda_2 = 0$ so we can do our analysis without the smoothing operator L, which is not approximated. Consider now the true solution:

$$\bar{x} = (A^T A + \lambda I)^{-1} A^T b \quad \text{(True Solution)}. \tag{3.13}$$

Notice that we can easily understand the significance of (3.13) by plugging in the (full rank) SVD $A = U \Sigma V^T$ into (3.13). One then obtains the solution:

$$\bar{x} = V D U^T b \quad \text{with} \quad D = \text{Diag}\left(\frac{\sigma_1}{\sigma_1^2 + \lambda}, \frac{\sigma_2}{\sigma_2^2 + \lambda}, \ldots, \frac{\sigma_{\text{r}}}{\sigma_{\text{r}}^2 + \lambda}, 0, \ldots, 0\right).$$

We see that the effect of the regularization is to filter the effects of singular vectors corresponding to small singular values σ_i, by replacing each σ_i by $\frac{\sigma_i}{\sigma_i^2 + \lambda}$, which prevents the singular vectors corresponding to singular values smaller than λ from dominating the solution [19].

Next, we state the approximate $\tilde{x_1}, \hat{x_1}, \tilde{x_2}, \tilde{x_3}$ that have been described in detail in the last section, but now with $\lambda_2 = 0$:

$$\tilde{x_1} = (A_k^T A_k + \lambda I)^{-1} A_k^T b, \tag{3.14}$$

$$\hat{x_1} = (A_k^T A_k + \lambda I)^{-1} A^T b, \tag{3.15}$$

$$\tilde{x_2} = \left((U_k^T A)^T (U_k^T A) + \lambda I\right)^{-1} (U_k^T A) U_k^T b, \tag{3.16}$$

$$\tilde{x_3} = V_k \left(\Sigma_k^2 + \lambda I\right)^{-1} \Sigma_k U_k^T b. \tag{3.17}$$

We have previously shown that $\tilde{x_2}$ and $\tilde{x_1}$ have the same solution. We will show in this section that $\tilde{x_3}$ also has the same solution as $\tilde{x_1}$. Notice that while the application of Tikhonov minimization acts to filter the small singular values of A on the solution, the use of the low rank SVD removes many of the small values entirely: the filtering is now done on those singular values which are retained.

We now introduce the Woodbury Inverse formula [23] which we will use in our analysis.

Lemma 3.5 *Take $D \in \mathbb{R}^{n\times n}$, $P \in \mathbb{R}^{n\times k}$, $T \in \mathbb{R}^{k\times k}$, and $R \in \mathbb{R}^{k\times n}$. Assume that D and T are invertible. Then $D + PTR$ is invertible if and only if $T^{-1} + RD^{-1}P$ is, and the following identity holds:*

$$(D + PTR)^{-1} = D^{-1} - D^{-1}P\left(T^{-1} + RD^{-1}P\right)^{-1}RD^{-1}. \tag{3.18}$$

Proof. The easiest proof of the identity is given using block matrix inversion.

$$\begin{pmatrix} D & P \\ R & -T^{-1} \end{pmatrix}\begin{pmatrix} X \\ Y \end{pmatrix} = \begin{pmatrix} I \\ 0 \end{pmatrix},$$

which reduces to:

$$DX + PY = I\left(\implies X = D^{-1}(I - PY)\right) \text{ and } RX - T^{-1}Y = 0\,(\implies Y = TRX)\,.$$

Plugging $Y = TRX$ into the first equation we get: $(D+PTR)X = I$ and plugging $X = D^{-1}(I-PY)$ into the second equation we get: $RD^{-1}(I - PY) = T^{-1}Y$. This can be expanded as:

$$\begin{aligned} RD^{-1} - RD^{-1}PY = T^{-1}Y \quad &\implies \quad RD^{-1} = (RD^{-1}P + T^{-1})Y \\ &\implies \quad (RD^{-1}P + T^{-1})^{-1}RD^{-1} = Y. \end{aligned}$$

We now substitute this Y into $DX + PY = I$ to get:

$$DX + P(RD^{-1}P + T^{-1})^{-1}RD^{-1} = I \implies X = D^{-1} - D^{-1}P(RD^{-1}P + T^{-1})^{-1}RD^{-1}.$$

But from $(D + PTR)X = I$ we have that $X = (D + PTR)^{-1}$ so we have the identity:

$$(D + PTR)^{-1} = D^{-1} - D^{-1}P(RD^{-1}P + T^{-1})^{-1}RD^{-1}.$$

□

Using the Woodbury inverse formula, we can derive expressions relating the terms $(A_k^TA_k + \lambda I)^{-1}$ and $(A^TA + \lambda I)^{-1}$ which appear in the solutions $\tilde{x_1}, \bar{x_1}, \tilde{x_2}, \tilde{x_3}$ and in the true solution $\bar{x}$.

Lemma 3.6 *Let k be in the range $1 \le k \le r-1$ and $\lambda > 0$. Then:*

$$\begin{aligned} &(A_k^TA_k + \lambda I)^{-1} = \lambda^{-1}I - V_kS_kV_k^T \\ &\textit{with} \quad S_k = \text{Diag}\left(\frac{\sigma_s^2}{\lambda^2 + \lambda\sigma_s^2}\right) \quad \textit{for} \quad s = 1, \ldots, k, \end{aligned} \tag{3.19}$$

and:

$$(A^TA+\lambda I)^{-1} = (A_k^TA_k+\lambda I)^{-1} - \hat{V}_k\hat{S}_k\hat{V}_k^T \tag{3.20}$$
$$\textit{with} \quad \hat{S}_k = \mathrm{Diag}\left(\frac{\sigma_s^2}{\lambda^2+\lambda\sigma_s^2}\right) \quad \textit{for} \quad s=k+1,\ldots,r.$$

These imply that:

$$\bar{x} = \left((A_k^TA_k+\lambda I)^{-1} - \hat{V}_k\hat{S}_k\hat{V}_k^T\right)A^Tb, \tag{3.21}$$
$$\tilde{x_1} = \left(\lambda^{-1}I - V_kS_kV_k^T\right)A_k^Tb, \tag{3.22}$$
$$\hat{x_1} = \left(\lambda^{-1}I - V_kS_kV_k^T\right)A^Tb. \tag{3.23}$$

Proof. The proof follows by the use of the Woodbury inverse formula:

$$(PTR+D)^{-1} = D^{-1} - D^{-1}P(RD^{-1}P+T^{-1})^{-1}RD^{-1}.$$

We match this with $(A_k^TA_k+\lambda I)^{-1} = (V_k\Sigma_k^2V_k^T+\lambda I)^{-1}$ to get $P=V_k$, $R=V_k^T$, $T=\Sigma_k^2$, and $D=\lambda I$:

$$\begin{aligned}
(A_k^TA_k+\lambda I)^{-1} &= \lambda^{-1}I - \lambda^{-1}V_k(V_k^T\lambda^{-1}V_k+\Sigma_k^{-2})^{-1}V_k^T\lambda^{-1} = \lambda^{-1}I - \lambda^{-2}V_k\left(\Sigma_k^{-2}+\lambda^{-1}I\right)^{-1}V_k^T \\
&= \lambda^{-1}I - \lambda^{-2}V_k\left(\mathrm{Diag}(\sigma_1^{-2},\ldots,\sigma_k^{-2})+\lambda^{-1}I\right)^{-1}V_k^T \\
&= \lambda^{-1}I - \lambda^{-2}V_k\,\mathrm{Diag}(\sigma_1^{-2}+\lambda^{-1},\ldots,\sigma_k^{-2}+\lambda^{-1})^{-1}V_k^T \\
&= \lambda^{-1}I - \lambda^{-2}V_k\,\mathrm{Diag}\left((\sigma_1^{-2}+\lambda^{-1})^{-1},\ldots,(\sigma_k^{-2}+\lambda^{-1})^{-1}\right)V_k^T \\
&= \lambda^{-1}I - \lambda^{-2}V_k\,\mathrm{Diag}\left(\frac{\lambda\sigma_1^2}{\lambda+\sigma_1^2},\ldots,\frac{\lambda\sigma_k^2}{\lambda+\sigma_k^2}\right)V_k^T \\
&= \lambda^{-1}I - V_k\,\mathrm{Diag}\left(\frac{\sigma_1^2}{\lambda^2+\lambda\sigma_1^2},\ldots,\frac{\sigma_k^2}{\lambda^2+\lambda\sigma_k^2}\right)V_k^T = \lambda^{-1}I - V_kS_kV_k^T,
\end{aligned}$$

which proves (3.19).

For (3.20), we have:

$$(A^TA+\lambda I)^{-1} = (A_k^TA_k+\hat{V}_k\hat{\Sigma}_k^2\hat{V}_k^T+\lambda I)^{-1} = (\hat{V}_k\hat{\Sigma}_k^2\hat{V}_k^T+M)^{-1},$$

with $M = A_k^TA_k+\lambda I$. Using Woodbury matrix formula:

$$(\hat{V}_k\hat{\Sigma}_k^2\hat{V}_k^T+M)^{-1} = M^{-1} - M^{-1}\hat{V}_k\left(\hat{\Sigma}_k^{-2}+\hat{V}_k^TM^{-1}\hat{V}_k\right)^{-1}\hat{V}_k^TM^{-1}.$$

Now, by (3.19) we have $M^{-1} = \lambda^{-1}I - V_kS_kV_k^T$ and by orthogonality we have $\hat{V}_k^TV_k = 0$:

$$\begin{aligned}
\hat{V}_k^TM^{-1} &= \hat{V}_k^T(\lambda^{-1}I - V_kS_kV_k^T) = \lambda^{-1}\hat{V}_k^T \\
M^{-1}\hat{V}_k &= (\lambda^{-1}I - V_kS_kV_k^T)\hat{V}_k = \lambda^{-1}\hat{V}_k.
\end{aligned}$$

Thus:

$$\begin{aligned}
(A^TA+\lambda I)^{-1} &= M^{-1} - M^{-1}\hat{V}_k\left(\hat{\Sigma}_k^{-2} + \hat{V}_k^T M^{-1}\hat{V}_k\right)^{-1}\hat{V}_k^T M^{-1} \\
&= M^{-1} - \lambda^{-1}\hat{V}_k\left(\hat{\Sigma}_k^{-2} + \hat{V}_k^T\lambda^{-1}\hat{V}_k\right)^{-1}\lambda^{-1}\hat{V}_k^T \\
&= M^{-1} - \lambda^{-2}\hat{V}_k\left(\hat{\Sigma}_k^{-2} + \lambda^{-1}I\right)^{-1}\hat{V}_k^T \\
&= M^{-1} - \lambda^{-2}\hat{V}_k\,\mathrm{Diag}\left(\frac{\lambda+\sigma_{k+1}^2}{\lambda\sigma_{k+1}^2},\ldots,\frac{\lambda+\sigma_r^2}{\lambda\sigma_r^2}\right)^{-1}\hat{V}_k^T \\
&= M^{-1} - \lambda^{-2}\hat{V}_k\,\mathrm{Diag}\left(\frac{\lambda\sigma_{k+1}^2}{\lambda+\sigma_{k+1}^2},\ldots,\frac{\lambda\sigma_r^2}{\lambda+\sigma_r^2}\right)\hat{V}_k^T \\
&= (A_k^TA_k+\lambda I)^{-1} - \hat{V}_k\,\mathrm{Diag}\left(\frac{\sigma_1^2}{\lambda^2+\lambda\sigma_{k+1}^2},\ldots,\frac{\sigma_r^2}{\lambda^2+\lambda\sigma_r^2}\right)\hat{V}_k^T \\
&= (A_k^TA_k+\lambda I)^{-1} - \hat{V}_k\hat{S}_k\hat{V}_k^T,
\end{aligned}$$

which proves (3.20).
Equations (3.19) and (3.20) imply that:

$$\begin{aligned}
\bar{x} &= (A^TA+\lambda I)^{-1}A^Tb = \left((A_k^TA_k+\lambda I)^{-1} - \hat{V}_k\hat{S}_k\hat{V}_k^T\right)A^Tb, \\
\tilde{x_1} &= (A_k^TA_k+\lambda I)^{-1}A_k^Tb = \left(\lambda^{-1}I - V_kS_kV_k^T\right)A_k^Tb, \\
\hat{x_1} &= (A_k^TA_k+\lambda I)^{-1}A^Tb = \left(\lambda^{-1}I - V_kS_kV_k^T\right)A^Tb.
\end{aligned}$$

□

Now we show that $\tilde{x_3}$ (involving the inversion of a $k\times k$ matrix) has the same solution as $\tilde{x_1}$ and derive the expression for the difference between $\tilde{x_1}$ and $\hat{x_1}$.

Lemma 3.7 *Let $\bar{x}$ be the solution of* (3.13)*, $\tilde{x_1}$ the solution of* (3.14)*, $\hat{x_1}$ the solution of* (3.15) *and $\tilde{x_3}$ the solution of* (3.17)*. Then, we have:*

$$\tilde{x_3} = \tilde{x_1}, \tag{3.24}$$

and

$$\hat{x_1} - \tilde{x_1} = \lambda^{-1}\left(A^T - A_k^T\right)b = \lambda^{-1}\hat{A_k}^T. \tag{3.25}$$

Proof. First note that:

$$V_kV_k^TA_k^Tb = V_kV_k^TV_k\Sigma_kU_k^Tb = V_k\Sigma_kU_k^Tb = A_k^Tb.$$

Next, we expand:

$$\begin{aligned}
\tilde{x_1} &= \left(\lambda^{-1}I - V_k S_k V_k^T\right) A_k^T b = \lambda^{-1} A_k^T b - V_k S_k V_k^T A_k^T b = \lambda^{-1} V_k V_k^T A_k^T b - V_k S_k V_k^T A_k^T b \\
&= V_k \left(\lambda^{-1}I - S_k\right) V_k^T A_k^T b = V_k \left(\lambda^{-1}I - \mathrm{Diag}\left(\frac{\sigma_s^2}{\lambda^2 + \lambda\sigma_s^2}\right)\right) V_k^T A_k^T b \\
&= V_k \left(\mathrm{Diag}\left(\frac{1}{\lambda} - \frac{\sigma_s^2}{\lambda^2 + \lambda\sigma_s^2}\right)\right) V_k^T A_k^T b = V_k \left(\mathrm{Diag}\left(\frac{(\sigma_s^2 + \lambda) - \sigma_s^2}{\lambda(\sigma_s^2 + \lambda)}\right)\right) V_k^T A_k^T b \\
&= V_k \,\mathrm{Diag}\left(\frac{1}{\sigma_s^2 + \lambda}\right) V_k^T A_k^T b = V_k (\Sigma_k^2 + \lambda I)^{-1} V_k^T A_k^T b = \tilde{x_3},
\end{aligned}$$

which proves (3.24). Next, for the difference between $\tilde{x_1}$ and $\hat{x_1}$ we have:

$$\begin{aligned}
\tilde{x_1} &= \left(A_k^T A_k + \lambda I\right)^{-1} A_k^T b = \left(\lambda^{-1}I - V_k S_k V_k^T\right) A_k^T b = \lambda^{-1} A_k^T b - V_k S_k V_k^T A_k^T b, \\
\hat{x_1} &= \left(A_k^T A_k + \lambda I\right)^{-1} A^T b = \left(\lambda^{-1}I - V_k S_k V_k^T\right) A^T b = \lambda^{-1} A^T b - V_k S_k V_k^T A^T b.
\end{aligned}$$

Note that:

$$V_k S_k V_k^T A^T b = V_k S_k V_k^T (V_k \Sigma_k U_k^T + \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T) b = V_k S_k V_k^T A_k^T b.$$

Hence:

$$\hat{x_1} - \tilde{x_1} = \lambda^{-1} A^T b - \lambda^{-1} A_k^T b = \lambda^{-1} (A^T - A_k^T) b = \lambda^{-1} \hat{A_k}^T b,$$

which proves (3.25). □

By the result of Lemma 3.7, the only solutions which differ from each other are $\tilde{x_1}$ and $\hat{x_1}$. We now analyze these two solutions with respect to the true solution $\bar{x}$.

Proposition 3.8 *Let $\bar{x}$ be the solution of* (3.13) *and $\tilde{x_1}$ the solution of* (3.14)*. Then:*

$$||\bar{x} - \tilde{x_1}||_2 \leq \frac{\sigma_{k+1}}{\lambda + \sigma_{k+1}^2} ||b||_2, \tag{3.26}$$

and

$$\tilde{x_1} = V_k V_k^T \bar{x}. \tag{3.27}$$

Proof. Recall that $\tilde{x_1} = (A_k^T A_k + \lambda I)^{-1} A_k^T b$ and that $\bar{x} = (A^T A + \lambda I)^{-1} A^T b$. Next by Lemma 3.6 and using that $A_k \hat{V}_k = (U_k \Sigma_k V_k^T) \hat{V}_k = 0$ and $\hat{V}_k^T V_k = 0$:

$$\begin{aligned}
(A^T A + \lambda I)^{-1} A^T &= (A^T A + \lambda I)^{-1} (A_k^T + \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T) = \left((A_k^T A_k + \lambda I)^{-1} - \hat{V}_k \hat{S}_k \hat{V}_k^T\right) (A_k^T + \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T) \\
&= (A_k^T A_k + \lambda I)^{-1} A_k^T + (A_k^T A_k + \lambda I)^{-1} \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T - \hat{V}_k \hat{S}_k \hat{\Sigma}_k \hat{U}_k^T \\
&= (A_k^T A_k + \lambda I)^{-1} A_k^T + (\lambda^{-1} I - V_k S_k V_k^T) \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T - \hat{V}_k \hat{S}_k \hat{\Sigma}_k \hat{U}_k^T \\
&= (A_k^T A_k + \lambda I)^{-1} A_k^T + \lambda^{-1} \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T - \hat{V}_k \hat{S}_k \hat{\Sigma}_k \hat{U}_k^T \\
&= (A_k^T A_k + \lambda I)^{-1} A_k^T + \hat{V}_k \left(\lambda^{-1} \hat{\Sigma}_k - \hat{S}_k \hat{\Sigma}_k\right) \hat{U}_k^T.
\end{aligned}$$

Since $\hat{S}_k = \text{Diag}(\frac{\sigma_s^2}{\lambda^2+\lambda\sigma_s^2})$ for $s = (k+1), \ldots, r$:

$$\lambda^{-1}\hat{\Sigma}_k - \hat{S}_k\hat{\Sigma}_k = \text{Diag}\left(\frac{\sigma_s}{\lambda} - \frac{\sigma_s^3}{\lambda(\lambda+\sigma_s^2)}\right) = \text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right) \quad \text{for} \quad s = (k+1), \ldots, r.$$

Hence:

$$(A^TA + \lambda I)^{-1}A^T = (A_k^TA_k + \lambda I)^{-1}A_k^T + \hat{V}_k\,\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\hat{U}_k^T,$$

which implies:

$$\begin{aligned}
\bar{x} &= (A^TA + \lambda I)^{-1}A^Tb = (A_k^TA_k + \lambda I)^{-1}A_k^Tb + \hat{V}_k\,\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\hat{U}_k^T && (3.28)\\
&= \tilde{x_1} + \hat{V}_k\,\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\hat{U}_k^Tb && (3.29)
\end{aligned}$$

$$\begin{aligned}
\implies ||\bar{x} - \tilde{x_1}||_2 &= \left\|\hat{V}_k\,\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\hat{U}_k^Tb\right\| = \left\|\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\hat{U}_k^Tb\right\| \leq \left\|\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\right\|_2 ||b||_2\\
&\leq \frac{\sigma_{k+1}}{\lambda+\sigma_{k+1}^2}||b||_2,
\end{aligned}$$

which proves (3.26).

Next, to derive (3.27), we have:

$$A^Tb = (V_k\Sigma_kU_k^T)b + (\hat{V}_k\hat{\Sigma}_k\hat{U}_k^T)b,$$

so that

$$\hat{V}_k\hat{\Sigma}_k\hat{U}_k^Tb = A^Tb - (V_k\Sigma_kU_k^T)b \implies \hat{V}_k^T\hat{V}_k\hat{\Sigma}_k\hat{U}_k^Tb = \hat{\Sigma}_k\hat{U}_k^Tb = \hat{V}_k^TA^Tb - 0 \implies \hat{U}_k^Tb = \hat{\Sigma}_k^{-1}\hat{V}_k^TA^Tb$$

$$\begin{aligned}
\implies \hat{U}_k^Tb &= \hat{\Sigma}_k^{-1}\hat{V}_k^T(A^TA + \lambda I)\bar{x} = \hat{\Sigma}_k^{-1}\hat{V}_k^T(V_k\Sigma_k^2V_k^T + \hat{V}_k\hat{\Sigma}_k^2\hat{V}_k^T + \lambda I)\bar{x} = \hat{\Sigma}_k^{-1}\left(\hat{\Sigma}_k^2\hat{V}_k^T + \lambda\hat{V}_k^T\right)\bar{x}\\
&= \hat{\Sigma}_k\hat{V}_k^T\bar{x} + \lambda\hat{\Sigma}_k^{-1}\hat{V}_k^T\bar{x} = (\hat{\Sigma}_k + \lambda\hat{\Sigma}_k^{-1})\hat{V}_k^T\bar{x}.
\end{aligned}$$

Using (3.28), we have:

$$\begin{aligned}
\bar{x} &= \tilde{x_1} + \hat{V}_k\,\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\hat{U}_k^Tb = \tilde{x_1} + \hat{V}_k\,\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)(\hat{\Sigma}_k + \lambda\hat{\Sigma}_k^{-1})\hat{V}_k^T\bar{x}\\
&= \tilde{x_1} + \hat{V}_k\,\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\text{Diag}\left(\sigma_s + \frac{\lambda}{\sigma_s}\right)\hat{V}_k^T\bar{x} = \tilde{x_1} + \hat{V}_k\,\text{Diag}\left(\frac{\sigma_s}{\lambda+\sigma_s^2}\right)\text{Diag}\left(\frac{\sigma_s^2+\lambda}{\sigma_s}\right)\hat{V}_k^T\bar{x}\\
&= \tilde{x_1} + \hat{V}_k\hat{V}_k^T\bar{x} = \tilde{x_1} + (I - V_kV_k^T)\bar{x} = \tilde{x_1} + \bar{x} - V_kV_k^T\bar{x}.
\end{aligned}$$

This proves (3.27):

$$\tilde{x_1} = V_kV_k^T\bar{x}.$$

□

Next, we look at the solution $\hat{x_1} = (A_k^T A_k + \lambda I)^{-1} A^T b$. Recall the difference here between $\tilde{x_1}$ is that $A^T b$ is not approximated by $A_k^T b$.

Proposition 3.9 *Let $\bar{x}$ be the solution of (3.13) and $\hat{x_1}$ the solution of (3.15). Then:*

$$||\bar{x} - \hat{x_1}||_2 \leq \frac{\sigma_{k+1}^3}{\lambda^2 + \lambda\sigma_{k+1}^2} ||b||_2, \tag{3.30}$$

and

$$\frac{||\bar{x} - \hat{x_1}||_2}{||\bar{x}||_2} \leq \frac{\sigma_{k+1}^2}{\lambda}. \tag{3.31}$$

Proof. We use lemma 3.6 to relate $\bar{x}$ to $\hat{x_1}$.

$$\begin{aligned} \bar{x} &= (A^T A + \lambda I)^{-1} A^T b = \left((A_k^T A_k + \lambda I)^{-1} - \hat{V}_k \hat{S}_k \hat{V}_k^T \right) A^T b & (3.32) \\ &= \hat{x_1} - \hat{V}_k \hat{S}_k \hat{V}_k^T A^T b = \hat{x_1} - \hat{V}_k \hat{S}_k \hat{V}_k^T \left(A_k^T + \hat{V}_k \hat{\Sigma}_k \hat{U}_k^T \right) b = \hat{x_1} - \hat{V}_k \hat{S}_k \hat{\Sigma}_k \hat{U}_k^T b, & (3.33) \end{aligned}$$

where the last equality follows from $\hat{V}_k^T A_k^T = 0$ and $\hat{V}_k^T \hat{V} = I$. Thus, we have:

$$||\bar{x} - \hat{x_1}||_2 = ||\hat{V}_k \hat{S}_k \hat{\Sigma}_k \hat{U}_k^T b||_2 = ||\hat{S}_k \hat{\Sigma}_k \hat{U}_k^T b||_2 \leq ||\hat{S}_k \hat{\Sigma}_k||_2 ||\hat{U}_k^T b||_2 \leq ||\hat{S}_k \hat{\Sigma}_k||_2 ||b||_2.$$

Now from lemma 3.6:

$$\begin{aligned} & \hat{S}_k \hat{\Sigma}_k = \text{Diag}\left(\frac{\sigma_{k+1}^3}{\lambda^2 + \lambda\sigma_{k+1}^2}, \frac{\sigma_{k+3}^2}{\lambda^2 + \lambda\sigma_{k+2}^2}, \ldots, \frac{\sigma_r^3}{\lambda^2 + \lambda\sigma_r^2} \right) \\ \implies \quad & ||\hat{S}_k \hat{\Sigma}_k||_2 = \max(\hat{S}_k \hat{\Sigma}_k) = \frac{\sigma_{k+1}^3}{\lambda^2 + \lambda\sigma_{k+1}^2} \\ \implies \quad & ||\hat{S}_k \hat{\Sigma}_k||_2 ||b||_2 = \frac{\sigma_{k+1}^3}{\lambda^2 + \lambda\sigma_{k+1}^2} ||b||_2. \end{aligned}$$

So we obtain the bound (3.30):

$$||\bar{x} - \hat{x_1}||_2 \leq ||\hat{S}_k \hat{\Sigma}_k||_2 ||b||_2 = \frac{\sigma_{k+1}^3}{\lambda^2 + \lambda\sigma_{k+1}^2} ||b||_2.$$

In order to obtain (3.31), we need to get rid of the $||b||_2$ term. We appeal back to (3.33):

$$\begin{aligned} \bar{x} &= \hat{x_1} - \hat{V}_k \hat{S}_k \hat{V}_k^T A^T b = \hat{x_1} - \hat{V}_k \hat{S}_k \hat{V}_k^T \left(A^T A + \lambda I \right) \bar{x} \\ &= \hat{x_1} - \hat{V}_k \hat{S}_k \hat{V}_k^T \left(V_k \Sigma_k^2 V_k^T + \hat{V}_k \hat{\Sigma}_k^2 \hat{V}_k^T + \lambda I \right) \bar{x} = \hat{x_1} - \hat{V}_k \hat{S}_k \left(\hat{\Sigma_k^2} + \lambda I \right) \hat{V}_k^T \bar{x}. \end{aligned}$$

It follows that:

$$\begin{aligned} & ||\bar{x} - \hat{x_1}||_2 = ||\hat{V}_k \hat{S}_k \left(\hat{\Sigma_k^2} + \lambda I \right) \hat{V}_k^T \bar{x}||_2 \leq ||\hat{V}_k \hat{S}_k \left(\hat{\Sigma_k^2} + \lambda I \right) \hat{V}_k^T||_2 ||\bar{x}||_2 = ||\hat{S}_k \left(\hat{\Sigma_k^2} + \lambda I \right)||_2 ||\bar{x}||_2 \\ \implies \quad & \frac{||\bar{x} - \hat{x_1}||_2}{||\bar{x}||_2} \leq ||\hat{S}_k \left(\hat{\Sigma_k^2} + \lambda I \right)||_2 \leq ||\hat{S}_k||_2 || \left(\hat{\Sigma_k^2} + \lambda I \right)||_2 = \frac{\sigma_{k+1}^2}{\lambda^2 + \lambda\sigma_{k+1}^2} (\sigma_{k+1}^2 + \lambda), \end{aligned}$$

which simplifies to:

$$\frac{||\bar{x} - \hat{x_1}||_2}{||\bar{x}||_2} \leq \frac{\sigma_{k+1}^2}{\lambda}.$$

□

Let us now recall some results we have derived. First of all, we have showed that $\tilde{x_1}$, $\tilde{x_2}$ and $\tilde{x_3}$ lead to the same solution. Numerically, however, one may still observe some differences if they are not run to convergence. On the other hand, $\tilde{x_1}$ and $\hat{x_1}$ differ from each other and have the following absolute error bounds with respect to the true solution $\bar{x}$:

$$\begin{aligned} ||\bar{x} - \tilde{x_1}||_2 &\leq \frac{\sigma_{k+1}}{\lambda + \sigma_{k+1}^2}||b||_2; \\ ||\bar{x} - \hat{x_1}||_2 &\leq \frac{\sigma_{k+1}^3}{\lambda\left(\lambda + \sigma_{k+1}^2\right)}||b||_2 \\ \implies \quad ||\bar{x} - \hat{x_1}||_2 &\leq \frac{\sigma_{k+1}^2}{\lambda}||\bar{x} - \tilde{x_1}||_2. \end{aligned}$$

Recall that the difference between the two is in the right hand side: $\hat{x_1}$ uses the un-approximated right hand side, or at least one computed with the wavelet transformed matrix (i.e. $A^T b \approx W^{-1} M^T b$). We mention again that one operation with a large A or M is not prohibitively expensive as it can be done by splitting the matrix into small enough blocks. For a large matrix, we would typically use $\lambda > 1$. In that case, based on the above relations, the solution $\hat{x_1}$ is closer to the true solution $\bar{x}$ when k is taken large enough so that the last singular value $\sigma_k < 1$. Thus, for a fixed k it may be better to use the approximation $\hat{x_1}$ when the matrix is significantly ill-conditioned, otherwise $\tilde{x_1}$ can be used. Another observation about the solution $\hat{x_1}$ compared to $\tilde{x_1}$ (and the other solutions equivalent to it) is that $\hat{x_1}$ for the same choice of λ is expected to have a larger norm:

Lemma 3.10 *Let $\tilde{x_1}$ be the solution of* (3.14) *and $\hat{x_1}$ the solution of* (3.15)*. Then, we have that $||\tilde{x_1}||_2 \leq ||\hat{x_1}||_2$.*

Proof. Recall that $A = A_k + \hat{A}_k$ and

$$\tilde{x_1} = (A_k^T A_k + \lambda I)^{-1} A_k^T b \quad ; \quad \hat{x_1} = (A_k^T A_k + \lambda I)^{-1} A^T b.$$

Now by Lemma 3.7:

$$\hat{x_1} = \tilde{x_1} + \lambda^{-1} \hat{A_k^T} b.$$

Thus, the norms are related as:

$$||\hat{x_1}||_2^2 = ||\tilde{x_1}||_2^2 + 2\lambda^{-1} \tilde{x_1}^T \hat{A_k^T} b + ||\hat{A_k^T} b||_2^2,$$

where the middle term is zero since $\hat{A}_k A_k^T = \hat{A}_k V_k = 0$:

$$\left(\tilde{x_1}^T \hat{A_k^T}\right)^T = \hat{A}_k \tilde{x_1} = \hat{A}_k (A_k^T A_k + \lambda I)^{-1} A_k^T b = \hat{A}_k (\lambda^{-1} I - V_k S_k V_k^T) A_k^T b = \lambda^{-1} \hat{A}_k A_k^T b + \hat{A}_k V_k S_k V_k^T A_k^T b = 0.$$

Thus:

$$||\hat{x_1}||_2^2 = ||\tilde{x_1}||_2^2 + ||\hat{A_k^T}b||_2^2 \implies ||\tilde{x_1}||_2 \leq ||\hat{x_1}||_2.$$

□

Thus, when using $\hat{x_1}$ as an estimate for $\bar{x}$ we typically would like to take a larger value of λ to obtain a solution with similar norm to that of $\tilde{x_1}$. If we use the same λ we will find that the components of the solution have larger amplitudes.

4 Numerical Experiments

In this section, we give some numerical examples to illustrate the approximation techniques we have discussed. We will use matrices from the Geotomography application which we have previously referred to. The full matrix A consists of 20 different blocks and has dimensions $2,968,933 \times 3,637,248$:

$$A = \begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_{20} \end{bmatrix}.$$

This matrix, in uncompressed form, is more than 3 TB in size and is thus too large for us to handle directly. In our illustrations, we will use also the smaller submatrix A_1 of the full matrix A. The submatrix has dimensions $438,674 \times 3,637,248$ and is about 115 GB in uncompressed form. We can load this matrix into memory. It is not hard to see why the matrix A is so large. In our application, the rows of the matrix correspond to earthquake-receiver pairs, the number of which is very high (almost 3 million). It is to our advantage to include as many such pairs as possible. The more rows we include, the more information we include in the system and the more detailed the solution. Likewise, the columns correspond to the coordinate system that is used. Each row is a sensitivity kernel [12], computed at each grid point of the Earth, from the core mantle boundary to the surface. We use a cubed sphere coordinate system [16] which divides the region within the Earth between the core-mantle boundary and the surface into 37 depth layers each divided laterally into 6 chunks subdivided into 128×128 voxels. This translates into approximately 3.6 million columns. For finer resolution, even more voxels (columns) need to be used. The matrix is sparse, having approximately 1 percent nonzeros. Let us look at the storage requirements for such a matrix. A typical sparse format stores the dimensions, the total number of nonzeros, the number of nonzerors in each row (or column), and the column (or row) indices of all the nonzeros, followed by the floating point values of all the nonzeros. We typically use integers to represent everything but the floating point values for which we use floats or doubles. The resulting binary file can easily be several terabytes in size when the dimensions and number of nonzeros are large.

We will use wavelet compression and low rank SVD in our tests. Thus, we will refer to the following

quantities corresponding to A and A_1:

$$M = \begin{bmatrix} M_1 \\ M_2 \\ \vdots \\ M_{20} \end{bmatrix} = \begin{bmatrix} \mathrm{Thr}(A_1 W^T) \\ \mathrm{Thr}(A_2 W^T) \\ \vdots \\ \mathrm{Thr}(A_{20} W^T) \end{bmatrix} \quad \text{and} \quad U_{1_k} \Sigma_{1_k} V_{1_k}^T \approx A_1 \; ; \; U_k \Sigma_k V_k^T \approx A.$$

Each row of the matrix A (and by extension A_1) is a sensitivity kernel that is defined over a cubedsphere coordinate system, in which the contents at the surface of a sphere of a given radius are projected onto six faces of a cube. Thus, each kernel has information for several different depth layers (corresponding to different radii from the core-mantle boundary to the Earth's surface). For more information on the inverse problem setup, see [17]. We will consider matrix vector operations with these matrices and the solutions:

$$(A_1^T A_1 + \lambda I)^{-1} A_1^T b \quad \text{and} \quad (A^T A + \lambda I)^{-1} A^T b \quad \text{and} \quad (A^T A + \lambda_1 I + \lambda_2 L^T L)^{-1} A^T b.$$

with L a Laplacian smoothing operator. The solutions represent seismic wave velocity corrections with respect to a spherically symmetric model. Thus, the solution has values at all latitudes and longitudes for each depth layer between the core-mantle boundary and the Earth's surface. When we plot the kernel or a solution, we usually make the plot at a single depth over all latitudes and longitudes. We can also make a vertical cross section through the different depths. As a consequence of the matrix data, the quality of resolution at different depths can significantly differ. The matrices we use come from a surface wave data set [21], such that only the top few depth layers near the surface carry nonzero information and even the bottom of these layers can already offer limited resolution. Thus the quality of approximations can vary somewhat for different depth layers. In order for the reader to have an idea of the data set we use, we present below some checkerboard reconstructions. We define x_{chk} to be a checkerboard grid, over the top few layers (near the surface). The result is plotted in Figure 2 using the depth profile (a cross-section plot showing the model representation over all depth layers) and corresponding cubed sphere representations at certain depths (we plot at each depth layer shown the projection onto the six cube faces). Then we form $b = Ax_{chk}$ and solve the regularized system $(A^T A + \lambda I)x = A^T b$ with $\lambda = 5$. We plot the solution x in Figure 3 using the same formats. We are unable to load the matrix A into memory, so we use the wavelet transformed and thresholded matrix $M = \mathrm{Thr}(AW^T)$ to approximate the matrix vector operations with A. We provide examples of wavelet compression in the next section. The comparison between x_{chk} and x gives us a summary of what the data set can pick up. In particular, we see from Figure 3 that the resolution is limited to layers near the surface, as expected. We mention also that the checkers used are about the size of what we can successfully resolve. Using smaller checkers doesn't lead to good reconstructions, even near the surface. The checkerboard example is important to keep in mind, as it shows the clear limitation of the data set: we are unable to resolve features at all depths, nor are we able to resolve very small features. Hence, in this case, we can safely use relatively high compression ratio approximation methods we have discussed (using fairly aggressive thresholding in wavelet based methods or small k relative to matrix dimension in SVD based schemes). Even though the solutions which result from these methods may not resolve some fine scale features in direct comparisons, it is important to keep in mind that these fine scale features cannot be realistically explained by the data we have available. This is true in many applications similar to ours.

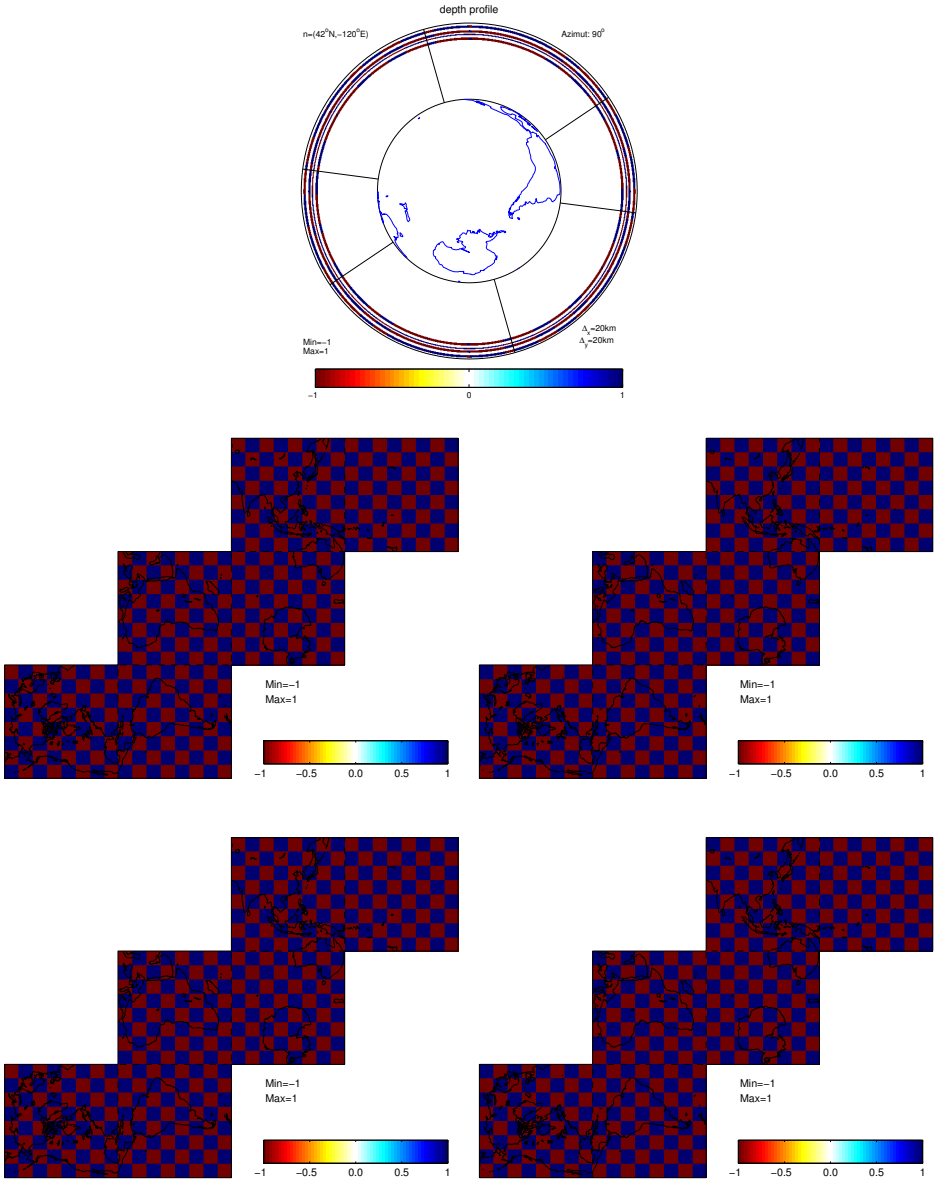

Figure 2: Checkerboard input (x_{chk}). Row 1: Depth profile. Row 2: layers 34 and 32 (135 and 316 km depth). Row 3: layers 30 and 28 (428 and 586 km depth).

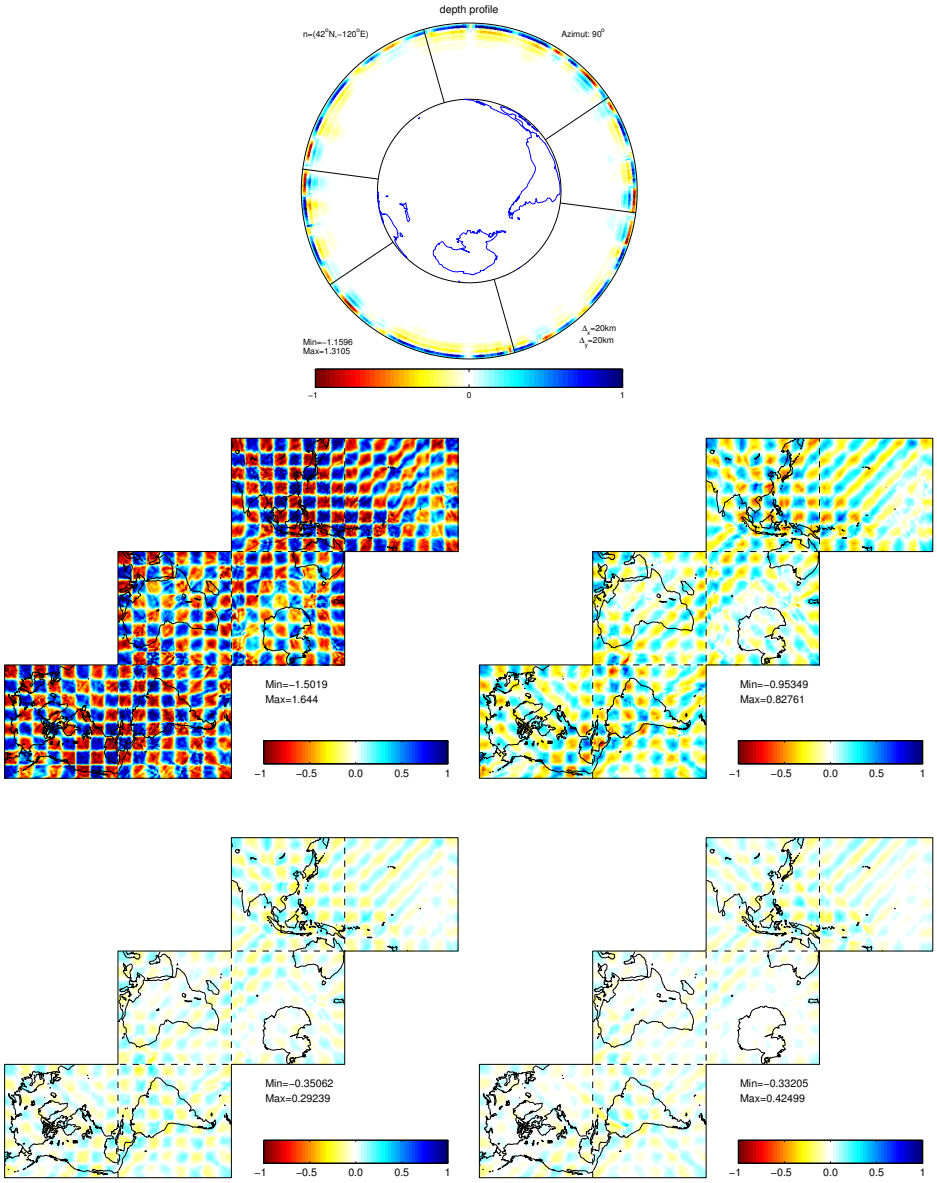

Figure 3: Checkerboard reconstruction (x). Row 1: Depth profile. Row 2: layers 34 and 32 (135 and 316 km depth). Row 3: layers 30 and 28 (428 and 586 km depth).

4.1 Wavelet Compression

We will illustrate some examples with wavelet compression using the matrices A_1 and $M_1 = \text{Thr}(A_1 W^T)$ since A_1 is small enough for us to load in uncompressed form. The size ratio between A_1 and M_1 depends on the compressibility of each row of A_1, which depends in turn on the implementation of the thresholding function $\text{Thr}(\cdot)$ and the wavelet transform W. Hence, we first illustrate the compression with a sample row of the matrix and then go on to compare matrix-vector operations with the whole matrix. We use the CDF $9-7$ wavelet transform for W. The thresholding is performed as follows: we sort the input vector by absolute magnitude and record the number of nonzeros. We then take as threshold the entry of the sorted array of nonzeros at the index defined by a fraction multiple of the number of nonzeros. All entries of the unsorted vector below this threshold are then set to zero. For this thresholding scheme, the lesser fraction multiple we choose, the more coefficients we will throw out and the worse the approximation but the better the compression ratio will be.

For each row of the matrix (or kernel) we proceed as follows. If we call the kernel row r, we take its wavelet transform and threshold, resulting in $\text{Thr}(Wr^T)$ which has significantly fewer nonzeros than r depending on the degree of thresholding and the wavelet transform used. To approximately reconstruct r from this compressed version, we use $\left(W^{-1}\,\text{Thr}(Wr)\right)^T$ and compare this to the original r. The reconstructed kernel is closer to the original for certain depths than for others. Below, in Figure 4, we plot a typical kernel r and its reconstructed version $\left(W^{-1}\left(\text{Thr}(Wr^T)\right)\right)^T$ as a function of percent cutoff at a certain depth near the surface. In the figure, the leftmost plot is the original kernel while the rightmost plot is the reconstructed kernel with about 10 percent of the coefficients retained after transforming. We see a notable degrade in quality. However, when we keep about 25 percent of the largest coefficients, we have much less noticeable reconstruction error. Notice that the error over all depths (all the entries of the kernel row) is greater than just at the particular depth layer at which it is plotted; but it is acceptable as long as we keep about 25 percent or more coefficients after transforming.

Next, we show what happens when we obtain the compressed matrix M_1 and use it to approximate matrix vector operations with A_1. Based on the plots in Figure 4, we retain about 30% of the coefficients in each row. Doing the transformation and thresholding for every row of A_1 gives us the matrix M_1. The full matrix A_1 is of size 115 GB while the matrix M_1 computed with our chosen threshold fraction comes out to be 35 GB. Figure 5 shows the percent errors for 50 randomly generated vectors between $A_1 x$ and $M_1 W^{-T} x$, $A_1^T y$ and $W^{-1} M_1^T y$, and between $A_1^T A_1 x$ and $W^{-1} M_1^T M_1 W^{-T} x$. We see that by approximating matrix-vector operations via the smaller matrix M_1 we introduce some errors. We see from Figure 5 that for this set of 50 random vectors that the errors in the approximate Ax and $A^T y$ operations are around 8% and that the errors in the approximate $A^T Ax$ operation sometimes reach 20% (although mostly hover closer to 10%). Iterative algorithms (like Conjugate Gradients) for solving the Tikhonov minimization system $(A_1^T A_1 + \lambda I)\bar{x} = A_1^T b$ rely mostly on the operation $A_1^T A_1 x$ so since the ℓ_2 functional is regularizing we expect similar errors to the final solution as we find in $A_1^T A_1 x$ (we do not expect these errors to blow up the error in the final solution).

Next, we discuss the use of these approximations in inversions. In Figure 6, we show the solutions from the regularization problems using the full matrix for matrix-vector operations compared with the

result obtained using the wavelet compressed matrix:

$$\begin{aligned} x_1 &= \arg\min_x \left\{ ||A_1 x - b_1||_2^2 + \lambda ||x||_2^2 \right\} \quad \text{using } A_1 \text{ for mat-vec mults;} \\ x_2 &= \arg\min_x \left\{ ||A_1 x - b_1||_2^2 + \lambda ||x||_2^2 \right\} \quad \text{using } M_1 = \text{Thr}(A_1 W^T) \text{ for mat-vec mults.} \end{aligned}$$

That is x_1 and x_2 solve the following linear systems:

$$(A_1^T A_1 + \lambda I) x_1 = A_1^T b \quad ; \quad (W^{-1} M_1^T M_1 W^{-T} + \lambda I) x_2 = W^{-1} M_1^T b.$$

To obtain x_2, we only keep M_1 in memory (which is about 3.5 times smaller than A_1). As clearly seen in Figure 6, after 300 iterations, the solutions x_1 and x_2 look very similar, although the error between x_1 and x_2 is about 10% at all depths and about 5% at the depth which is shown. Note that the percent error in the ℓ_2 norm between the solution which we report may not be a good measure of how similar the two solutions are since many differences are in small details below the scale which can be successfully resolved. Also in Figure 6 you can see the plots of solution norm and χ^2 value versus iteration for the two solutions. The curves for both are very close to each other. The norm of the solution is just the ℓ_2 norm of the iterate x^n at iteration n. The χ^2 value is calculated using the formula:

$$\chi^2 = \frac{1}{P} \sum_{k \text{ not outlier}} |r_k^n|^2,$$

where $r^n = Ax^n - b$ and $P = m - m_0$ (number of rows minus number of outliers). For each datum, we define standard errors (estimated a priori), then scale the system to be univariant (i.e. all standard errors are equal to 1). We define outliers as entries of the vector r^n that are not within three standard errors. In the inversions we present, the outliers are identified after 5 and 25 iterations, corresponding to dips in the χ^2 that may be seen in the plots. Since our systems are univariant, we would like for the χ^2 of the converged solution to be close to one. However, this is not possible for this data set without including extra correction terms in the data. From the figures below we can make the following general observations about the wavelet approximation approach:

- This approach allows us to use a matrix about 3.5 times smaller (in this case). Greater compression ratios are possible with more advanced transforms tailored to data and coordinate system and/or quality tradeoffs.
- In matrix vector operations, the error in the approximation to operation $A^T Ax$ is greater than for the approximations to operations Ax and $A^T y$.
- The approximate solution obtained with a matrix is very close to the original.
- The norm and χ^2 values of the solution to the regularization problem are accurately computed when using the wavelet compressed matrix.

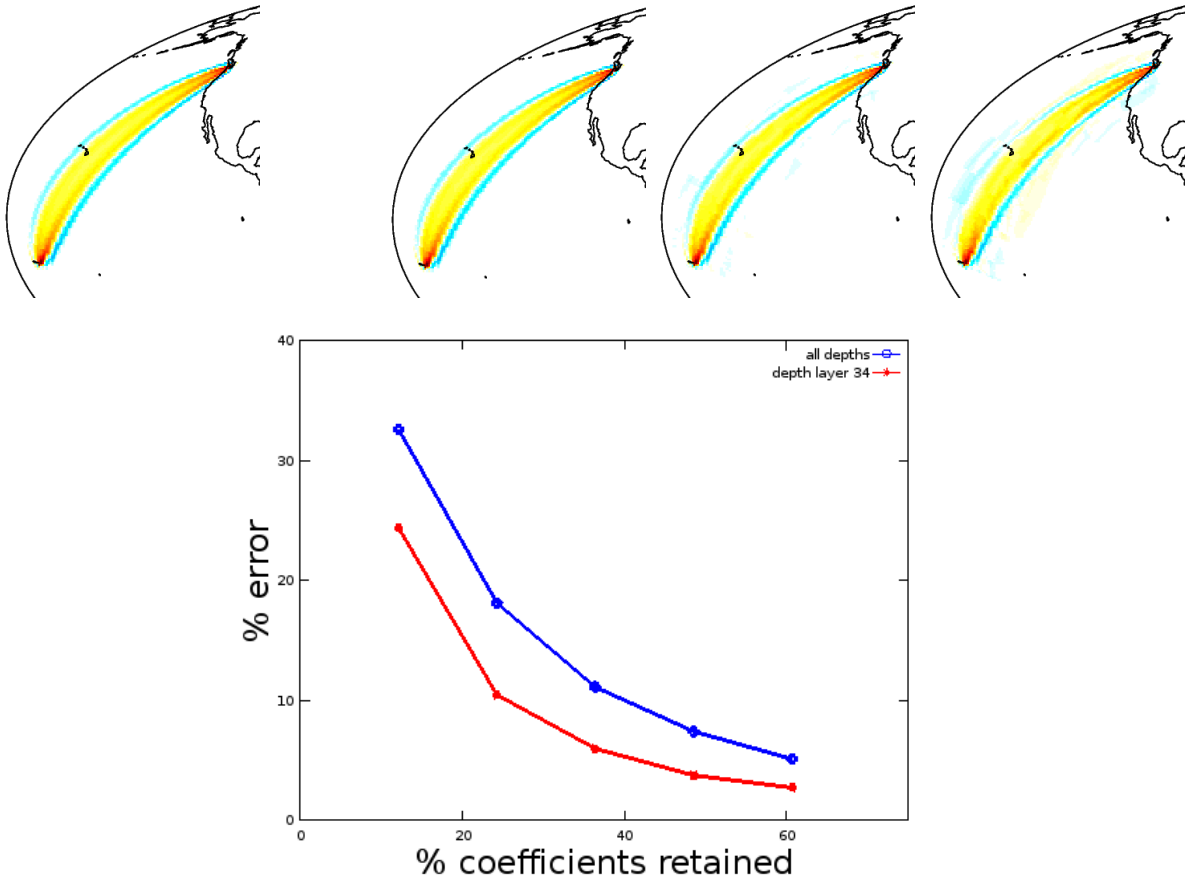

Figure 4: Original kernel and compressed kernels (at 135 km depth) with different numbers of coefficients retained after thresholding: approximately 48, 24, and 10 percent coefficients, respectively. The bottom left plot shows the percent error curve between the reconstructed and original kernel versus the number of nonzeros retained: errors for all depths and only for the displayed depth are shown.

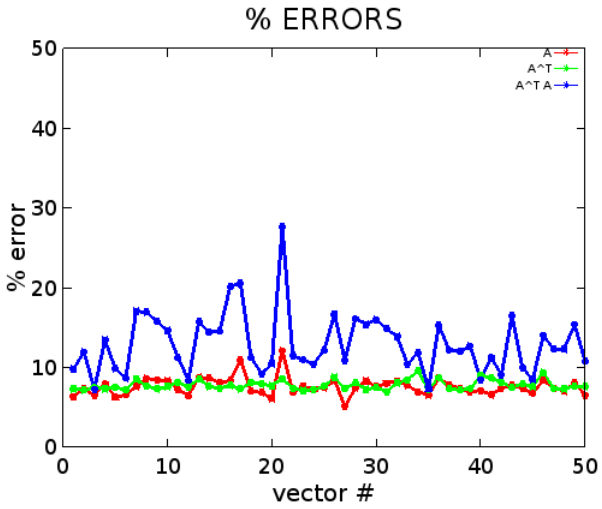

Figure 5: Percent errors between A_1x and $M_1W^{-T}x$, A_1^Ty and $W^{-1}M_1^Ty$, and $A_1^TA_1x$ and $W^{-1}M_1^TM_1W^{-T}x$ for 50 random vectors x and y.

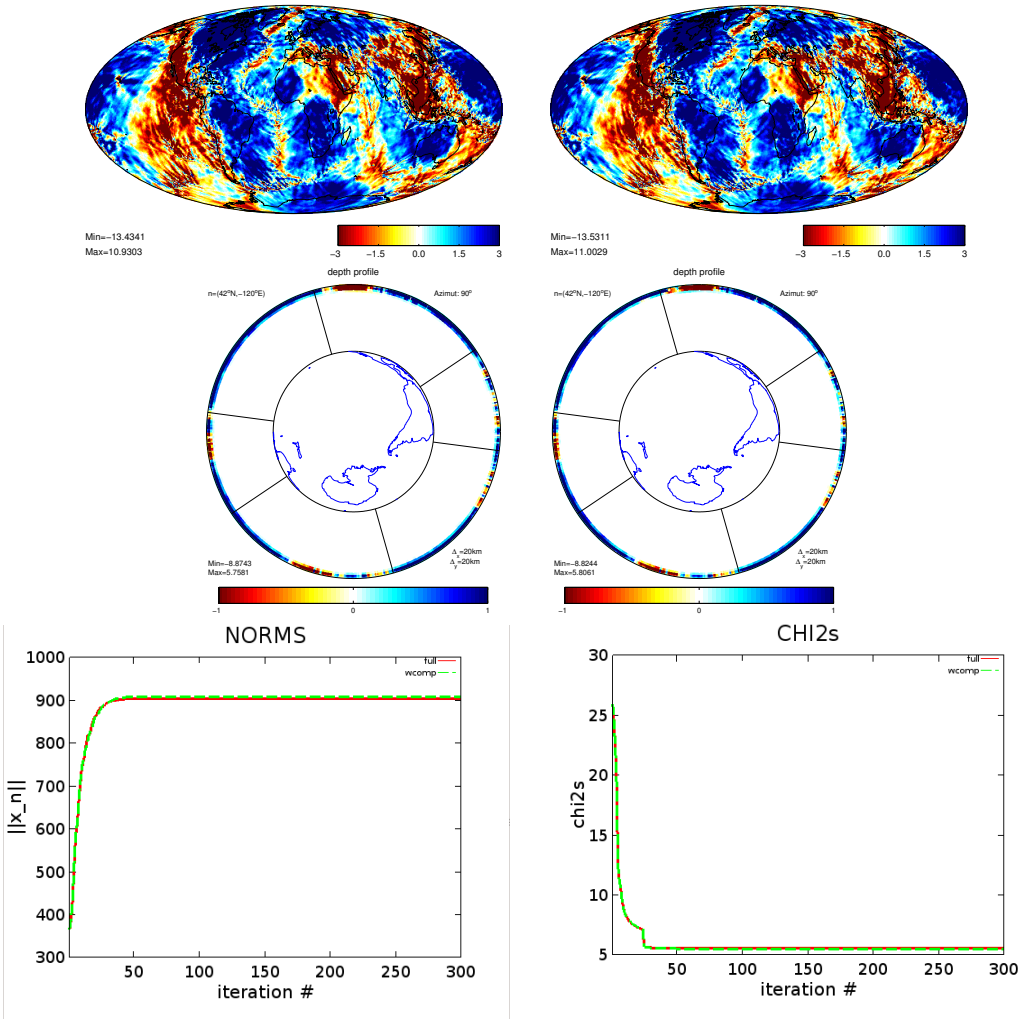

Figure 6: Regularized solutions x_1 and x_2 obtained using matrix-vector operations with full (left) and compressed (right) matrices. Row 1: plots of solutions at 135 km depth. Row 2: depth profiles. Row 3: solution norms and χ^2 values versus iteration.

4.2 Low Rank SVD Compression

We now illustrate some results of low rank SVD compression. We will use the first 2000 eigenvectors (that is, $k = 2000$), with the decomposition computed using the randomized algorithm previously discussed. We show results for both the smaller matrix A_1 with dimensions $(438,674 \times 3,637,248)$ and the larger matrix A with dimensions $(2,968,933 \times 3,637,248)$ containing A_1 as one of its 20 blocks. We compute the approximate low rank SVDs:

$$A_1 \approx U_{1_k} \Sigma_{1_k} V_{1_k}^T \quad \text{and} \quad A \approx U_k \Sigma_k V_k^T$$

using matrix vector operations approximated through the wavelet compressed matrices M_1 and M. We now state the sizes and dimensions of the matrices involved:

- A_1, dimensions $(438,674 \times 3,637,248)$, size is 115 GB
- M_1, dimensions $(438,674 \times 3,637,248)$, size is 35 GB
- U_{1_k}, Σ_{1_k}, V_{1_k}, dimensions $(438,674 \times 2000), (2000 \times 2000), (3637248 \times 2000)$, sizes are 7 GB, 30 MB, 55 GB ($\approx$ 62 GB total)
- A, dimensions $(2,968,933 \times 3,637,248)$, size is 3.2 TB (approximate, never computed)
- M, dimensions $(2,968,933 \times 3,637,248)$, size is 1 TB
- U_k, Σ_k, V_k, dimensions $(2,968,933 \times 2000), (2000 \times 2000), (3,637,248 \times 2000)$, sizes are 45 GB, 30 MB, 55 GB ($\approx$ 100 GB total)

Notice that for the matrix A_1, the total size of the SVD components (which are not sparse matrices) is greater than the size of the matrix M_1. Hence, it may not be practical to use the low rank SVD decomposition for this matrix, unless one is interested in using the $k \times k$ approach discussed in the previous section.

We now describe the figures that we present, which are similar to those presented for the wavelet compression approach. First in Figure 7, we illustrate the errors in matrix-vector operations which result from the SVD approximations. We compare results to the wavelet compressed approximate computations since the matrix A is not available in full form. For the smaller matrix A_1, we compare the operations $M_1 W^{-T} x \approx A_1 x$ to $U_{1_k} \Sigma_{1_k} V_{1_k}^T x$, $W^{-1} M_1^T y \approx A_1^T y$ to $V_{1_k} \Sigma_{1_k} U_{1_k}^T y$, and $W^{-1} M_1^T M_1 W^{-T} x \approx A_1^T A_1 x$ to $V_{1_k} \Sigma_{1_k}^2 V_{1_k}^T x$ and for the big matrix A, the same operations: $MW^{-T} x \approx Ax$ to $U_k \Sigma_k V_k^T x$, $W^{-1} M^T y \approx A^T y$ to $V_k \Sigma_k U_k^T y$, and $W^{-1} M^T M W^{-T} x \approx A^T A x$ to $V_k \Sigma_k^2 V_k^T x$. We illustrate the results for 50 randomly generated vectors in each case. Note that since A_1 and A are not well conditioned, $A_1^T A_1$ and $A^T A$ are much worser conditioned (the square of the condition number) and can thus be better approximated with the same rank approximation. This is exactly what we observe: the errors between $A_1^T A_1 x$ and $V_{1_k} \Sigma_{1_k}^2 V_{1_k}^T x$ are lower than between the other two operations. The same holds true when using the big matrix: for $A^T A$ versus A. This is opposite to what we observed in the case of wavelet compression where approximate operations with A_1 were slightly more accurate than with $A_1^T A_1$ and this has implications for calculations of quantities such as χ^2 which we will discuss. In

Figure 8, we illustrate the errors between the full matrix A_1 and the SVD quantities, which are similar to the errors between the SVD approximated operations and the compressed wavelet approximated ons. We see, as expected, that corresponding to the full matrix, the errors are slightly higher.

Next, we present the solutions obtained using the low rank SVD. We will show solutions to the following systems corresponding to $A_1 x = b_1$ and $Ax = b$:

$$
\begin{aligned}
&(W^{-1}M_1^T M_1 W^{-T} + \lambda I)x_2 = W^{-1}M_1^T b_1 \quad \text{wavelet compressed solution for } A_1 \\
&(V_{k_1}\Sigma_{k_1}^2 V_{k_1}^T + \lambda I)x_3 = V_{k_1}\Sigma_{k_1}U_{k_1}^T b_1 \quad \text{replacing all instances of } A_1 \text{ by low rank SVD} \\
&(V_{k_1}\Sigma_{k_1}^2 V_{k_1}^T + 5\lambda I)x_4 = W^{-1}M_1^T b_1 \quad \text{using the low rank SVD only on the left hand side} \\
&(W^{-1}M^T M W^{-T} + \lambda I)x_5 = W^{-1}M^T b \quad \text{wavelet compressed solution for } A \\
&(V_k\Sigma_k^2 V_k^T + \lambda I)x_6 = V_k\Sigma_k U_k^T b \quad \text{replacing all instances of } A \text{ by low rank SVD} \\
&(V_k\Sigma_k^2 V_k^T + 10\lambda I)x_7 = W^{-1}M^T b \quad \text{using the low rank SVD only on the left hand side}
\end{aligned}
$$

with $\lambda = 1$. The different solutions after 250 and 150 iterations are illustrated in Figures 9 and 10. As before, we plot the different solutions at a certain depth, plot the depth profiles at all depths, and show plots for the solution norms and χ^2 values versus iteration. We also show the following solutions corresponding to the system with Laplacian smoothing included:

$$
\begin{aligned}
&(W^{-1}M^T M W^{-T} + \lambda_1 I + \lambda_2 L^T L)x_8 = W^{-1}M^T b \quad \text{wavelet compressed with smoothing} \\
&(V_k\Sigma_k^2 V_k^T + \lambda_1 I + \lambda_2 L^T L)x_9 = V_k\Sigma_k U_k^T b \quad \text{SVD 1 with smoothing} \\
&(V_k\Sigma_k^2 V_k^T + 10\lambda_1 I + \lambda_2 L^T L)x_{10} = W^{-1}M^T b \quad \text{SVD 2 with smoothing}
\end{aligned}
$$

with $\lambda_1 = \lambda_2 = 1$. These solutions are illustrated in Figure 11. In all these systems, we use a higher λ when using the un-approximated right hand side based on the norm relations we showed in Lemma 3.10 so that all solutions have similar magnitudes.

We now discuss the different solutions. In Figure 9, we compare the wavelet compressed solution of x_2 (which was shown in the last section to be very close to x_1) to the two SVD solutions x_3 and x_4. From the plots we see that the solutions with the low rank SVD recover most features of the full solution and capture the behavior (but not the true value) of changing solution norm $||x_n||_2$ and χ^2 values. We do observe, however, that the approximate solution x_3 has significantly smaller norm than the solution x_2 and significantly higher χ^2 values. Both are somewhat expected: using $U_{1_k}\Sigma_{1_k}V_{1_k}^T$ instead of A_1 amounts to regularization in itself, completely removing the effects of singular vectors corresponding to small singular values smaller than σ_{k_1} and in effect adding to the damping imposed by Tikhonov regularization. The χ^2 values approximated with the matrix $U_{1_k}\Sigma_{1_k}V_{1_k}^T$ are simply not accurate because the matrix vector operations with this matrix have high errors (which we may observe in Figure 7) and so the residual vector r^n is not accurately computed. On the other hand, the solution is accurate since it involves mostly the $A_1^T A_1 x$ operation which is accurately approximated by $V_{1_k}\Sigma_{1_k}^2 V_{1_k}^T$. That is to say, in truth, the solution computed has lower χ^2 values than shown. This is clearly shown in the bar plot in Figure 9, where we compare the mean χ^2 value after 20 iterations from the two SVD solutions x_3 and x_4 computed using the low rank matrix $U_{1_k}\Sigma_{1_k}V_{1_k}^T$ and using the

wavelet compressed MW^{-T} approximation. The same solutions have correspondingly lower χ^2 values when the residual $r^n = A_1 x^n - b_1$ is approximated via $MW^{-T}x^n - b_1$ instead of $U_{1_k}\Sigma_{1_k}V_{1_k}^T x^n - b_1$.

In Figure 10 we show the three solutions corresponding to the full system $Ax = b$. Here the solution x_5 is quite detailed and we see that the low rank SVD solutions x_6 and x_7 miss many small details present in the wavelet compressed solution. We clearly see from the plot of solution norms versus iteration that the low rank SVD solutions converge faster than the solution with the wavelet compressed matrix. As we have previously discussed, this is expected since replacing the full matrix by the low rank SVD amounts to regularization and doing Tikhonov minimization with the low rank matrix now filters the effects of the singular vectors corresponding to the small singular values which are retained. We see also in this case that the solution x_7 reveals more details than x_6. However, for solutions x_4 and x_7 the threshold must be adjusted to recover similar magnitudes to the true solution. For the solutions with Laplacian smoothing shown in Figure 11, we can see that the low rank SVD solutions provide very good approximation; in this case, since the finer details of the solution are smoothed out by the Laplacian operator, the SVD solutions are even closer to the wavelet compressed solution.

Based on the figures discussed here, we can make the following conclusions:

- This approach allows one to use much smaller matrices (in the case of A, the low rank SVD components are collectively less than 30 times the size of the full matrix).
- However, one may experience premature convergence and loss of details (depending on the value of k used in the low rank SVD). In our examples, we used a very small k relative to the dimensions of matrix A. The approximate solutions have visible deviations in fine details from the wavelet compressed or original solutions. However, these fine details may not be properly resolved by the data set in the first place, given the resolution limitations of the data set. Using higher k would give more details, but requires greater computation and storage.
- In matrix vector operations, the error in the approximation to operation $A^T Ax$ is significantly less than for the approximations to operations Ax and $A^T y$.
- The χ^2 is not accurately computed if the low rank SVD matrix is used to compute the residual; instead one should use the wavelet compressed or full matrix (for one computation) to accurately estimate the χ^2 value of the solution vector.
- When the Laplacian penalty is used, the low rank SVD solutions provide an even better approximation since the fine details of the original solution are smoothed out.

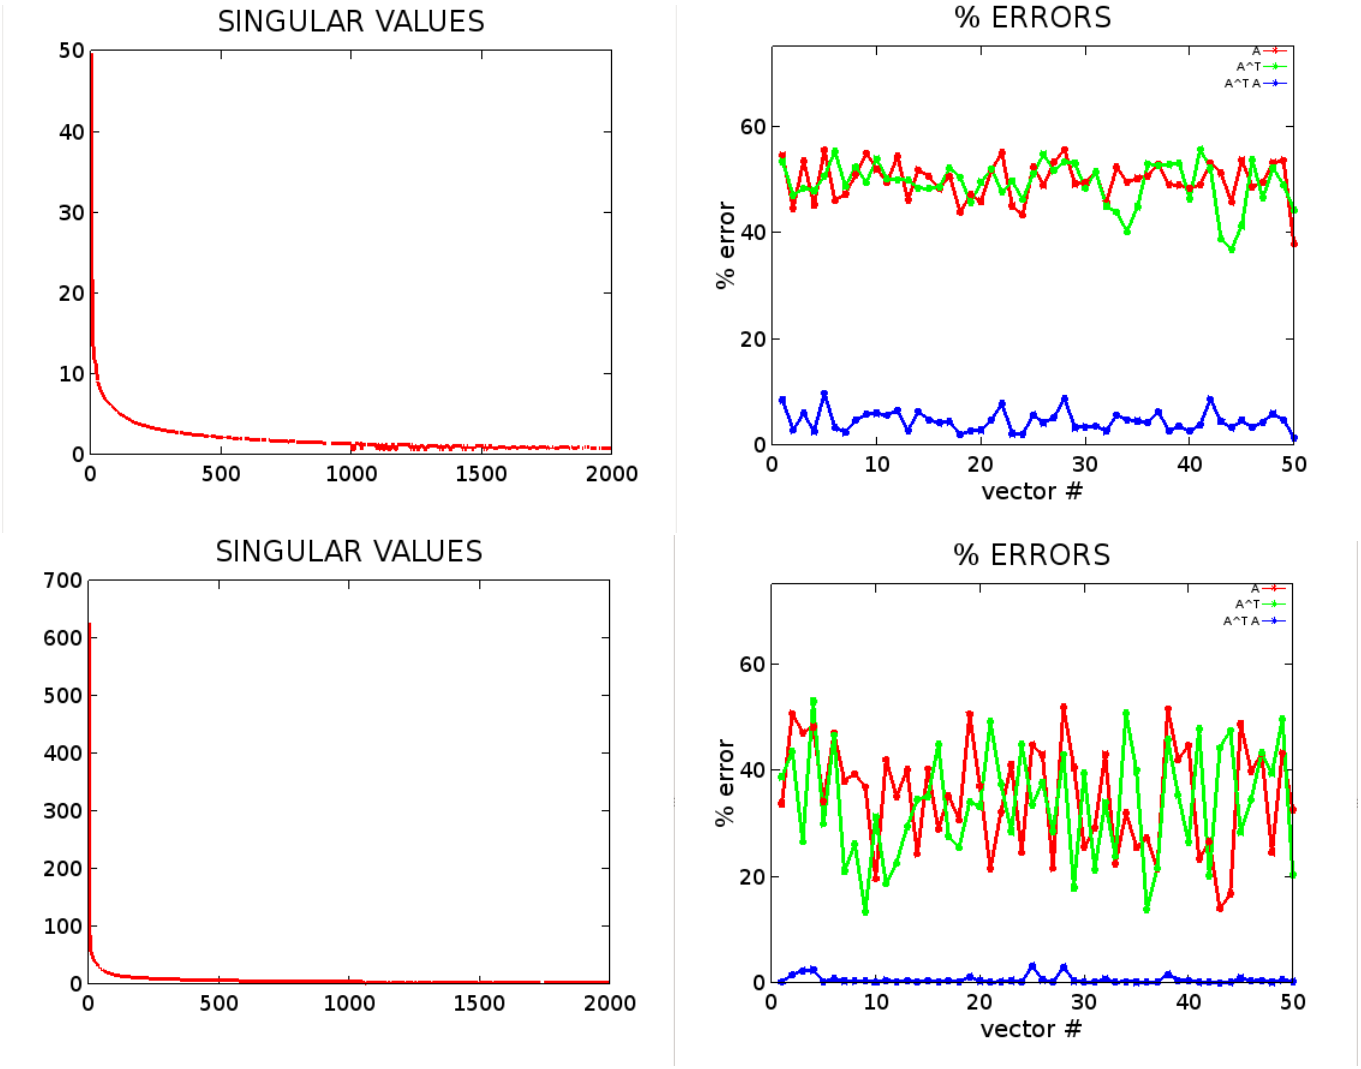

Figure 7: Row 1: the first 2000 calculated singular values of A_1 and percent errors with matrix-vector operations using low rank SVD compared to using the wavelet compressed operations via the matrix M_1. Row 2: the first 2000 calculated singular values of the big matrix A and percent errors with matrix-vector operations using low rank SVD compared to using the wavelet compressed matrix via the matrix M.

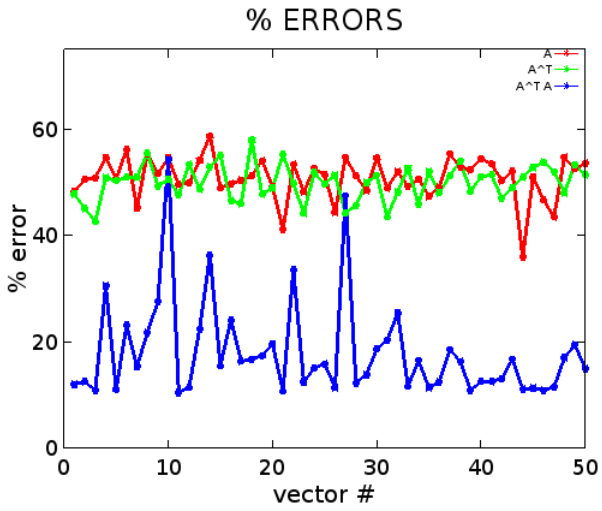

Figure 8: Percent errors in matrix-vector operations using low rank SVD compared to using full matrix A_1.

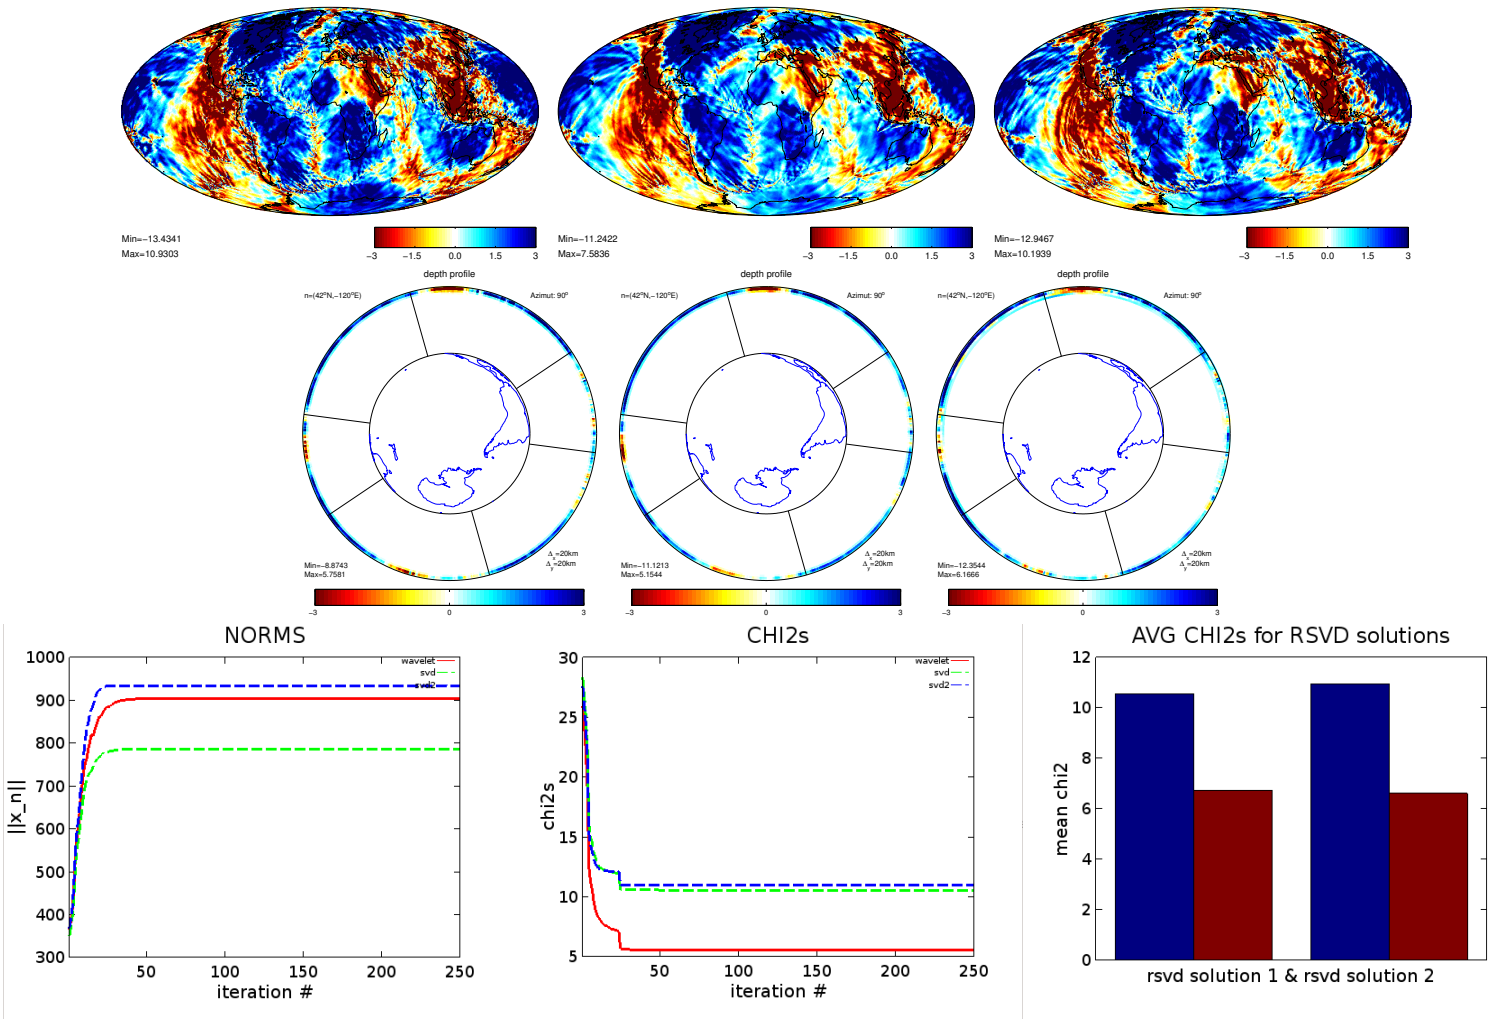

Figure 9: Plots for regularized solutions x_2, x_3, and x_4. First row: solutions plotted at 135 km depth. Second row: depth profiles for the solutions. Third row: norms of solution and χ^2 value at each iteration, bar plot of average χ^2 of the two SVD solutions computed with the low rank SVD matrix and the wavelet compressed matrix.

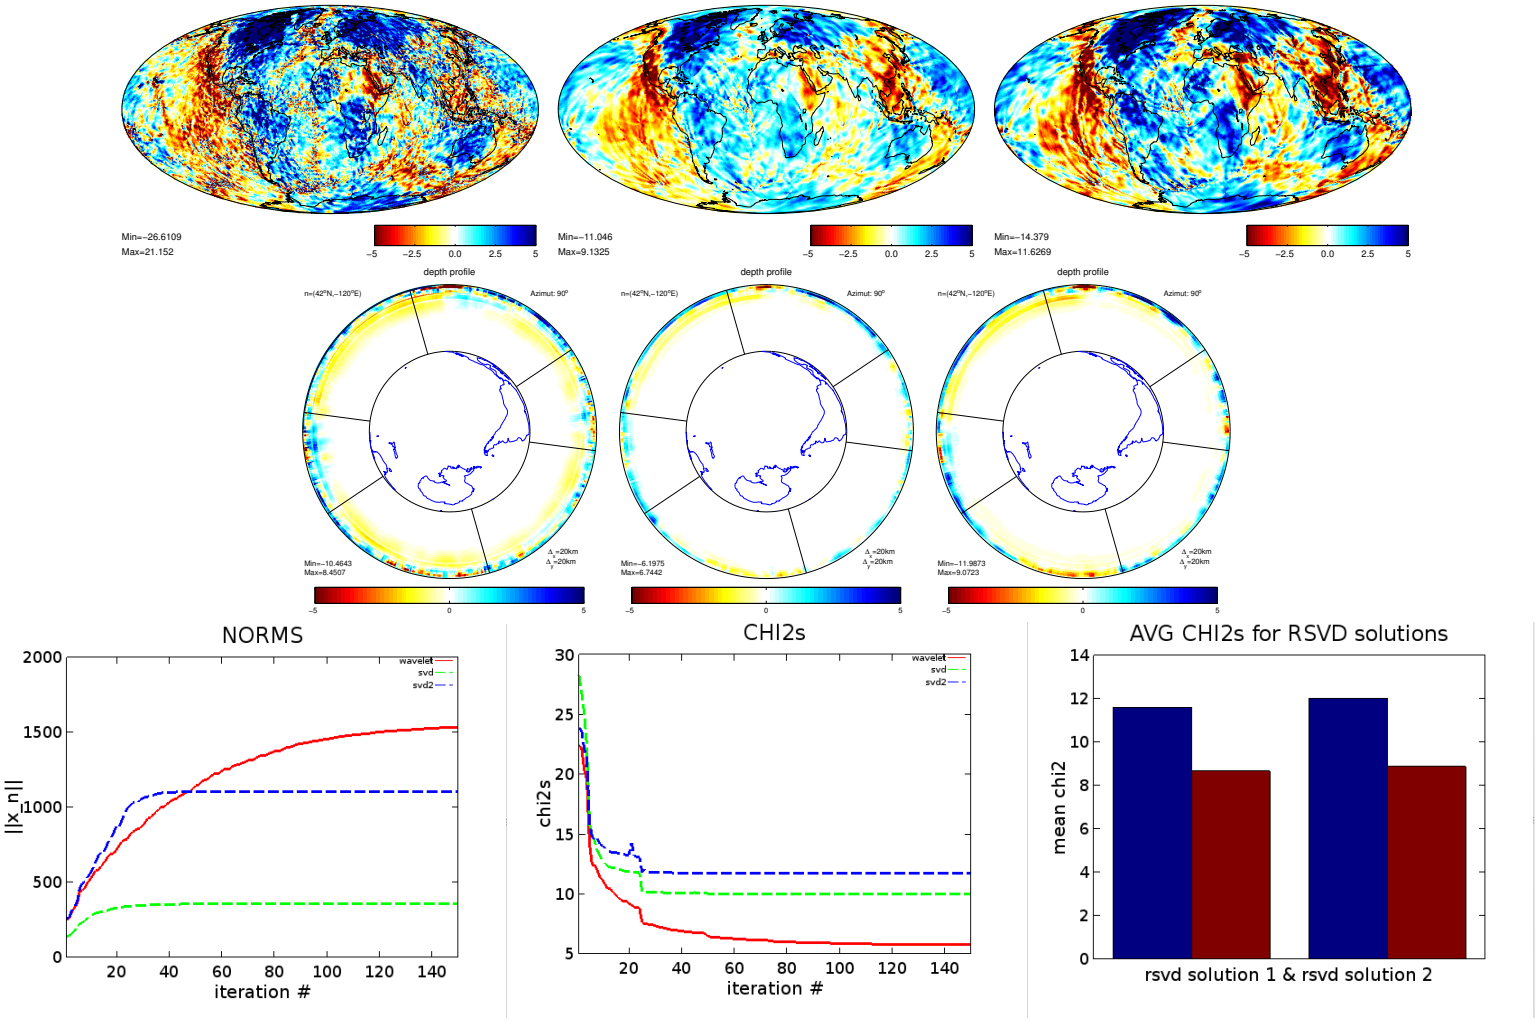

Figure 10: Plots for regularized solutions x_5, x_6, and x_7. First row: solutions plotted at 135 km depth. Second row: depth profiles for the solutions. Third row: norms of solution and χ^2 value at each iteration, bar plot of average χ^2 of the two SVD solutions computed with the low rank SVD matrix and the wavelet compressed matrix.

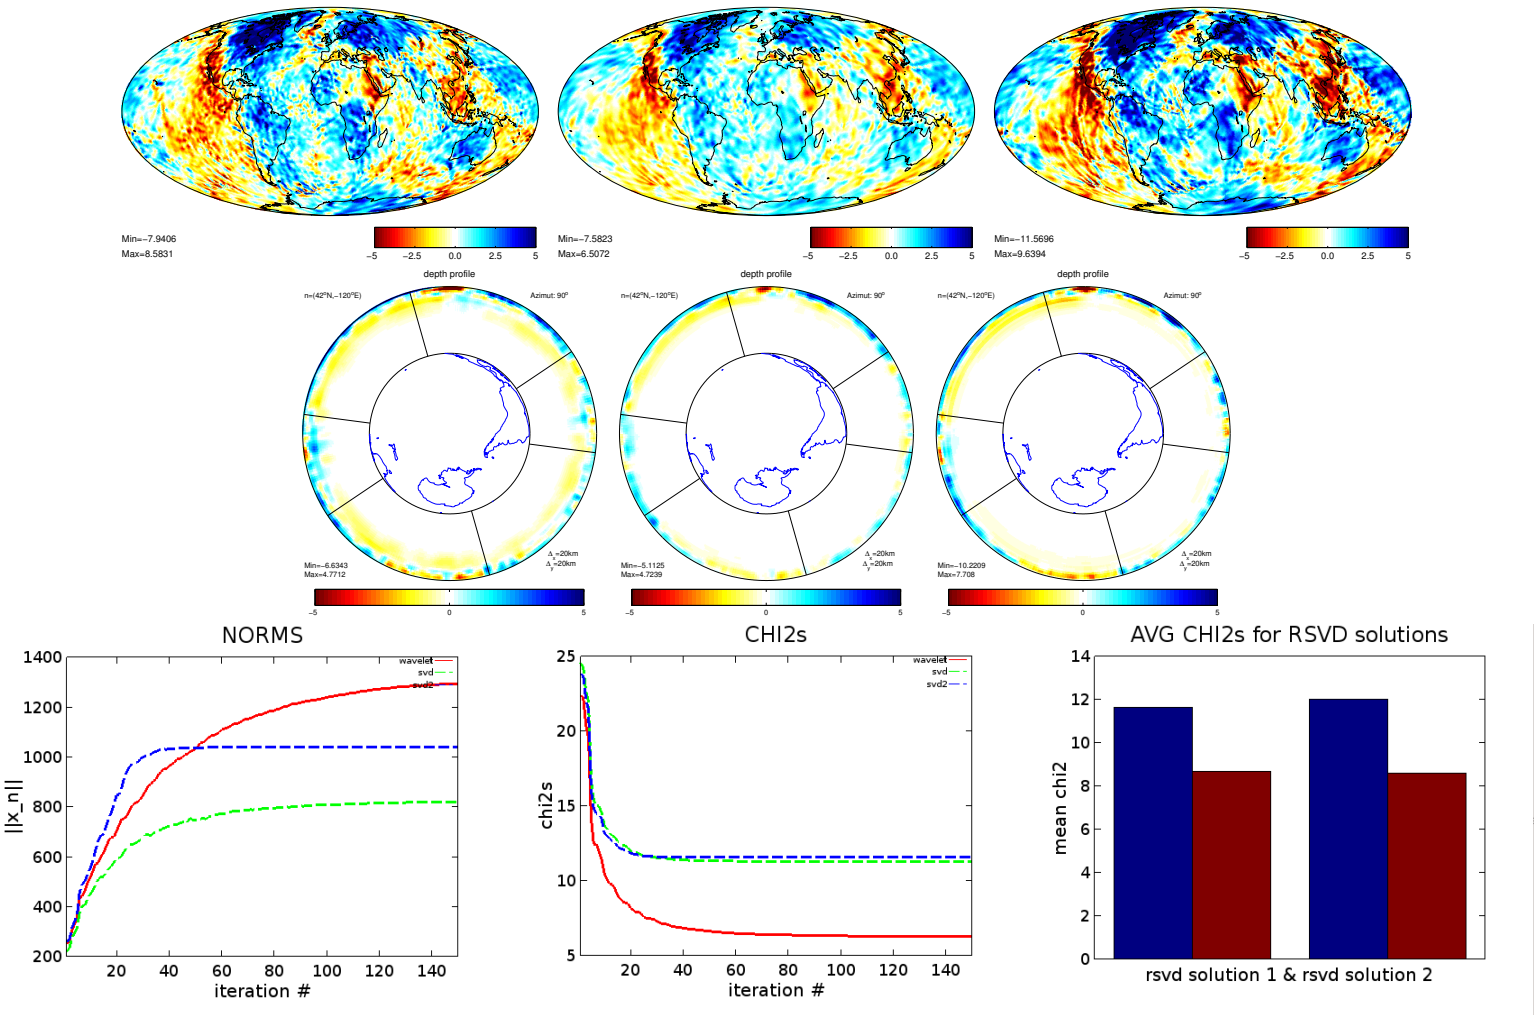

Figure 11: Plots for regularized solutions with Laplacian smoothing x_8, x_9, and x_{10}. First row: solutions plotted at 135 km depth. Second row: depth profiles for the solutions. Third row: norms of solution and χ^2 value at each iteration, bar plot of average χ^2 of the two SVD solutions computed with the low rank SVD matrix and the wavelet compressed matrix.

5 Conclusions

We have presented two different approaches for obtaining approximate solutions to regularization problems: using wavelet compression techniques and the low rank SVD. We illustrate the application of these techniques to ℓ_2 regularization for a large scale inverse problem from Geotomography. We have also presented some mathematical analysis for the various SVD based schemes we have considered, showing interesting equivalence between seemingly different schemes, with very different memory requirements. The techniques we present are also well applicable to other types of optimization problems. In fact, the methods presented here can be of use to any application where matrix-vector operations with large matrices are required, especially if the matrices are not well conditioned. The wavelet compressed approach is found to be very accurate and gives close reconstructions to the true solution, assuming the data is wavelet compressible. Based on the type of transform and data utilized a compression ratio of 3 or more can be expected. For large ill-conditioned matrices, the low rank SVD approach gives significantly better compression ratios (> 10) and resolves the main solution features. In our experiments, we note that the low rank SVD approximation is best for the matrix $A^T A$, because A is not ill-conditioned enough for the small rank k we choose, relative to the matrix dimensions. Hence, the estimates of various quantities such as χ^2 of the solution may not be accurate when the low rank SVD is used with small k. However, in the case of the regularization algorithms we consider, the solutions produced with the low rank SVD are of good quality since they depend mostly on matrix vector operations with matrix $A^T A$ which are well approximated even with small k. For both wavelet compressed and low rank SVD based methods, the accuracy and compression ratio are inversely proportional and controlled by the user. In the case of wavelet compression, the time it takes to form the compressed matrix is nearly independent of the threshold used. However, for the computation of the low rank SVD, the work involved substantially grows as the rank k increases. The ultimate choice of k or the wavelet threshold depends on the compression and accuracy requirements dictated by the user, as well as the resolving power of the data set. Often, a checkerboard style test can be done on a representative portion of the data to see the resolution a data set is capable of, in order to have an idea for the approximation accuracy required.

6 Acknowledgements

The authors would like to thank Ignace Loris, Gunnar Martinsson and Frederik Simons for very helpful advice and discussion on the paper content.

References

[1] A.N. Akansu and R.A. Haddad. *Multiresolution Signal Decomposition: Transforms, Subbands, and Wavelets*. Academic Press, Inc., Orlando, FL, USA, 1992.

[2] D. Calvetti, S. Morigi, L. Reichel, and F. Sgallari. Tikhonov regularization and the l-curve for large discrete ill-posed problems. *Journal of Computational and Applied Mathematics*, 123, 2000. Numerical Analysis 2000. Vol. III: Linear Algebra.

[3] J. Chárlety, S. Voronin, Nolet G., I. Loris, F.J. Simons, K. Sigloch, and I.C. Daubechies. Global seismic tomography with sparsity constraints: Comparison with smoothing and damping regularization. *Journal of Geophysical Research - Solid Earth*, 2013.

[4] A. Cohen, I.C. Daubechies, and J.-C. Feauveau. Biorthogonal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 45(5):485–560, 1992.

[5] I.C. Daubechies. Orthonormal bases of compactly supported wavelets. *Communications in Pure and Applied Mathematics*, 41:909–996, 1988.

[6] I.C. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11):1413–1457, 2004.

[7] E. Debayle and M. Sambridge. Inversion of massive surface wave data sets: Model construction and resolution assessment. *Journal of Geophysical Research*, 109:B02316, 2004.

[8] N. Halko, P.G. Martinsson, and J.A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.*, 53(2):217–288, May 2011.

[9] W. Härdle, G. Kerkyacharian, D. Picard, and A. Tsybokov. *Wavelets, Approximation, and Statistical Applications*, volume 129 of *Lecture Notes in Statistics*. Springer-Verlag, New York, 1998.

[10] J. Lampe, L. Reichel, and H. Voss. Large-scale tikhonov regularization via reduction by orthogonal projection. *Linear Algebra and its Applications*, 436(8):2845 – 2865, 2012. Special Issue dedicated to Danny Sorensen's 65th birthday.

[11] I. Markovsky. *Low Rank Approximation: Algorithms, Implementation, Applications*. Communications and Control Engineering. Springer, 2012.

[12] H. Marquering, G. Nolet, and F.A. Dahlen. Three-dimensional waveform sensitivity kernels. *Geophysical Journal International*, 132(3):521–534, 1998.

[13] Y. Meyer. *Wavelets: Algorithms & Applications*. Society for Industrial and Applied Mathematics, Philadelphia, 1993. Translated and revised by Robert D. Ryan.

[14] G. Nolet. *A Breviary of Seismic Tomography*. Cambridge Univ. Press, Cambridge, U.K., 2008.

[15] C.C. Paige and M.A. Saunders. Lsqr: An algorithm for sparse linear equations and sparse least squares. *ACM Trans. Math. Softw.*, 8(1):43–71, March 1982.

[16] C. Ronchi, R. Iacono, and P.S. Paolucci. The cubed sphere: A new method for the solution of partial differential equations in spherical geometry. *Journal of Computational Physics*, 124(1):93 – 114, 1996.

[17] F.J. Simons, I. Loris, G. Nolet, I.C. Daubechies, S. Voronin, J. S. Judd, P.A. Vetter, J. Chárlety, and C. Vonesch. Solving or resolving global tomographic models with spherical wavelets, and the scale and sparsity of seismic heterogeneity. *Geophysical Journal International*, 187(2):969–988, 2011.

[18] W. Sweldens. The lifting scheme: A new philosophy in biorthogonal wavelet constructions. In Andrew F. Laine, Michael A. Unser, and Mladen V. Wickerhauser, editors, *Wavelet applications in signal and image processing III*, volume 2569 of *Proceedings of SPIE*, pages 68–79, 1995.

[19] A.N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Math. Dokl*, 1963.

[20] L.N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, 1997.

[21] H.-J. van Heijst and J.H. Woodhouse. Global high-resolution phase velocity distributions of overtone and fundamental mode surface waves determined by mode branch stripping. *Geophysical Journal International*, 137(3):601–620, 1999.

[22] S. Wang and Z. Zhang. Improving cur matrix decomposition and the nystrom approximation via adaptive sampling. *J. Mach. Learn. Res.*, 14(1):2729–2769, January 2013.

[23] M.A. Woodbury. *Inverting modified matrices*. Statistical Research Group, Memo. Rep. no. 42. Princeton University, Princeton, N. J., 1950.